FINDING A MAXIMUM-GENUS GRAPH IMBEDDING

. Merrick L. Furst, Jonathan L. Gross, & Lyle A. McGeoch

Finding a Maximum-Genus Graph Imbedding

December 2, 1985

Merrick L. Furst [1]   *Jonathan L. Gross* [2]   *Lyle A. McGeoch* [1]

0. Abstract

The computational complexity of constructing the imbeddings of a given graph into surfaces of different genus is not well-understood. In this paper, topological methods and a reduction to linear matroid parity are used to develop a polynomial-time algorithm to find a maximum-genus cellular imbedding. This seems to be the first imbedding algorithm for which the running time is not exponential in the genus of the imbedding surface.

1. Introduction

Lower-dimensional topology has long been approached ·combinatorially. For most questions about imbeddings, there exist exhaustive algorithms. Since the number of combinatorial equivalence classes of graph imbeddings is a super-exponential function of the number of vertices, such exhaustive algorithms are computationally infeasible.

There have been several algorithmic achievements. Hopcroft and Tarjan [IIT] obtained a linear-time algorithm to test planarity of graphs, while Gross and Rosen [GR] showed how to test planarity of 2-complexes. Filotti [F] found a polynomial-time algorithm to determine if a cubic graph can be imbedded in the torus, and Filotti, Miller, and Reif [FMR] generalized this work with an algorithm to imbed a graph in a surface of minimum genus G in time $O(v^{O(G)})$. All these algorithms are based on extending partially imbedded graphs, and they all produce an

---

imbedding whenever it exists. Reif [R79] showed that there are limits to this approach by proving that the problem of deciding whether a partial imbedding in some surface can be extended to a full imbedding in that surface is NP-complete.

Our present concern is the determination of the "maximum genus" of a graph. There is no limit to the number of handles one might add to a surface in which a graph is already imbedded. For the concept of maximum genus to be meaningful, one must stipulate that every region of the imbedding be cellular – that is, the interior of the region must be homeomorphic to an open disk. Such a restriction is in no way artificial. It corresponds to restricting handles to be "essential".

Maximum-genus imbeddings, and the related notion of upper-imbeddable graphs, have received considerable attention in recent years. A graph is called *upper-imbeddable* if it has a maximum-genus imbedding with one or two faces. Nordhaus, Stewart, and White [NSW], Ringeisen [R70,R72], and Zaks [Z] showed that various classes of graphs were upper-imbeddable. Nebeský [N] and Jungerman [J] described combinatorial invariants of upper-imbeddable graphs. Xuong [X79b] showed that all graphs with two disjoint spanning trees, such as 4-edge connected graphs, are upper-imbeddable.

We consider the computational complexity of obtaining a maximum-genus imbedding. Our starting point is the combinatorial characterization by Xuong [X79a] of the maximum genus of a graph. This involves the consideration of all spanning trees of a graph, of which there can be exponentially many. We improve the obvious exponential-time algorithm to a polynomial-time algorithm.

2. Preliminaries about Topological Graph Theory

In topological graph theory, a "graph" is defined to be a (possibly) non-simplicial 1-complex. In other words, multiple adjacencies and self-loops are permitted. There are many reasons for this generality. In particular, the most powerful presently-known way to construct an imbedding of a large simplicial graph into a large-genus surface is to derive it as a branched covering of an imbedding of a smaller, non-simplicial graph — ideally with one vertex and many self-loops — into a smaller-genus surface. (See Gross and Tucker [GT] or White [W].)

In this paper, we consider only *simplicial (simple)* graphs: those without self-loops or multiple adjacencies. Any graph containing self-loops and multiple adjacencies can be transformed into a simplicial graph by inserting one or more vertices in the interior of these edges. Moreover, the resulting graph is homeomorphic to the original graph, and accordingly, it has the same maximum genus. This enables us to simplify the notation.

2.1. Graph terminology An *(undirected simplicial) graph* $G = (V, E)$ has a finite set of *vertices* V and a finite set of *edges* E, and an edge may be represented as an unordered pair of vertices (v, w), i.e. its endpoints. An edge is said to be *incident* on its endpoints. Two distinct edges are said to be *adjacent* if they are incident on a common vertex.

The *degree* of a vertex is the number of edges incident on it. A *path* from vertex v to vertex w is a sequence of edges (v, u_1), (u_1, u_2), ..., (u_k, w) in E, such that the vertices $v, u_1, u_2, ..., u_k$, and w are distinct. However, if the starting point v is the same as the final point w, the sequence is called a *cycle*. A graph is *connected* if there is a path between every pair of its vertices. A connected, acyclic graph is called a *tree*. By a *spanning tree* of graph G, we mean a subgraph that is a tree and contains every vertex.

The notation $G + e$ is an abbreviation for the graph $(V, E \cup \{e\})$, and the notation $G - e$ is an abbreviation for the graph $(V, E - \{e\})$.

A graph is *directed* if each edge is thought to have a beginning and an end. We represent a directed edge as an ordered pair. Unless it is otherwise obvious from context, the graphs discussed will be connected and undirected.

2.2. Surfaces Our terminology is compatible with that of Gross and Tucker [GT] and of White [W].

The *topological spaces* of interest here are all homeomorphic to subspaces of E^3. A *homeomorphism* between two topological spaces is a continuous bijective mapping with a continuous inverse. A connected topological space is a *surface* if every point has a neighborhood that is homeomorphic to the closed unit disk. A surface S is *orientable* if it does not contain a Möbius band.

We deal only with closed orientable surfaces. Every such surface S is homeomorphic to a generalized torus. The number of handles is denoted $\gamma(S)$ and is called the *genus* of the surface. A sphere, for example, is a surface of genus 0, a torus is a surface of genus 1, and a 2-handled torus is a surface of genus 2.

2.3. Graph imbeddings and faces Although a graph is an abstract combinatorial object, there is a topological representation of it: in Euclidean 3-space, we represent each vertex by a distinct point and each edge by a distinct curve between the two endpoints, where a *curve* means a homeomorphic image of the unit interval [0,1]. We require that the interior of an edge intersect no other edge or vertex of the graph. When referring to a graph in a topological setting, we mean such a representation.

An *imbedding* $G \to S$ of a graph G in the surface S is a continuous one-to-one mapping. The components of $S - G$ are called *regions*. If each region is homeomorphic to an open disk, the imbedding is *cellular*, and the regions are called *faces*. All our imbeddings are cellular. The set of faces of an imbedding is denoted F.

A *maximum-genus imbedding* of a connected graph is a cellular imbedding of the graph in an orientable surface having maximum genus among all such imbedding surfaces. The Euler polyhedral equation

$$|V| - |E| + |F| = 2 - 2\gamma(S)$$

holds for all cellular imbeddings. Thus, a maximum-genus imbedding is the same thing as a minimum-facecount imbedding.

2.4. Rotation systems A *rotation* at a vertex v is a cyclic permutation of the edges incident on it. Since our graphs are simplicial, we may specify a rotation at v in the format

$$v. \quad u_1 u_2 \dots u_d$$

where the vertices $u_1, \dots, u_d$ are the opposite endpoints of the edges incident on v. It follows that a vertex v with degree d admits $(d - 1)!$ different rotations.

A list of rotations, one for each vertex, is call a *rotation system* for the graph. This concept is due to Heffter [H]. Starting with a graph imbedding in an oriented surface, there corresponds an obvious rotation system, namely, the one in which the rotation at each vertex is consistent with the cyclic order of the neighboring vertices in that imbedding.

Edmonds [E] was first to call attention explicitly to a method for inverting that correspondence. To each oriented edge (u, v), one assigns the oriented edge (v, w) such that vertex w is the immediate successor of vertex u in the rotation at vertex v. The result is a permutation on the set of oriented edges, that is, on the set in which each undirected edge appears twice, once with each possible direction. In each edge-orbit under this permutation, the consecutive oriented edges line up head to tail, from which it follows that they form a directed cycle in the graph. We observe that it is possible for both orientations of the same edge to appear twice in the same edge-orbit. If there are n oriented edges in the orbit, then an n-sided polygon can be fitted into it. Fitting a polygon to every such edge-orbit results in a polygon on both sides of each edge, and collectively the polygons form a surface in which the graph is cellularly imbedded.

Sometimes one describes the rotation system of a graph pictorially, as in Figure 2.1. The graph is drawn in the plane so that the incidence of edges at each vertex is consistent with the rotation

4

system. Obviously, unless the rotation system happens to correspond to a planar imbedding, there will be edge-crossings in the drawing. Such a drawing permits one to trace along the edge-orbits, as illustrated. Since the graph shown has 6 vertices and 10 edges, and since the rotation system has 2 edge-orbits, the imbedding surface has Euler characteristic $6 - 10 + 2$, which equals -2, from which it follows that the imbedding surface has genus two.


Figure 2.1. A graph with two edge-orbits in its rotation system.

The existence of the bijective correspondence between the cellular imbeddings of a graph and the rotation systems enables us to reformulate the problem of finding the maximum genus of the graph as a problem of finding a rotation system with the minimum number of edge-orbits. Since edge-orbits correspond to boundary-walks of faces, this is equivalent to seeking a minimum-facecount imbedding.

We can depict the boundary-walk of each face of an imbedding as a directed graph with one directed edge for each traversal of the underlying undirected edge, and multiple copies of each vertex; the boundary-walks for the rotation system of Figure 2.1 are shown in Figure 2.2. Any closed boundary-walk can also be written as an alternating (cyclic) sequence of vertices and edges $v_1 e_1 v_2 e_2 \ldots e_k v_1$. A subsequence $e_i v e_{i+1}$ of a walk is called a *corner*, corresponding in an obvious way to the geometric corner of a face of a polygonal imbedding.

2.5. Adding and deleting edges If an edge is added to, or deleted from, an imbedded graph, then all faces in the imbedding are unchanged except those incident on that edge. Furthermore, either two faces are merged or one face is split into two faces.

5

Figure 2.2. The boundary-walks of the imbedding in Figure 2.1.

Suppose that an edge $e = (u, v)$ is added to a graph and its imbedding, so that its ends are inserted between two corners of one face. If the boundary-walk around the original face was of the form $v\alpha_1 w\alpha_2 v$, where α_i is a subwalk, then as illustrated by Figure 2.3, the new edge splits the old boundary-walk into two walks: $v\alpha_1 we_1 v$ and $w\alpha_2 ve_2 w$. Similarly, if an edge e that is common to two faces is deleted from an imbedding, then two boundary-walks are merged and the new imbedding has one less face.


Figure 2.3. Adding an edge across a face.

If an edge is added to a graph and its ends are inserted between corners of two different faces, then both those faces are merged into one larger face. In particular, suppose that new edge e runs from the corner of v in boundary-walk $v\alpha_1 v$ to the corner of w in boundary-walk $w\alpha_2 w$. Then a merged face results, with boundary-walk $ve_1 w\alpha_2 we_2 \alpha_1 v$, as depicted in Figure 2.4. Likewise, the deletion of an edge e occurring twice on one boundary-walk splits the corresponding face into two

Figure 2.4. Adding an edge between two faces.

smaller faces.

3. Maximum-Genus Imbeddings

We now direct our attention to the problem of constructing a maximum-genus imbedding. Xuong [X79a] proved that calculating the maximum genus of a graph is reducible to calculating the value of a combinatorial invariant which he called its *deficiency*.

The *deficiency* $\xi(G, T)$ *of a spanning tree* T in a graph G is defined to be the number of connected components of $G - T$ that contain an odd number of edges. The *deficiency* $\xi(G)$ *of a graph* G is defined to be the minimum tree deficiency over all spanning trees T of G. We call a spanning tree that realizes $\xi(G)$ a *Xuong tree*. Figure 3.1 shows a graph and one of its Xuong trees. Since the complement of the Xuong tree has two odd components, it follows that the graph has deficiency two.


Figure 3.1. A spanning tree (solid edges) with minimum deficiency.

The edge complement $G - T$ of any tree T is called a *cotree*. Tree T is a spanning tree if and only if $G - T$ is a minimum cotree. The number of edges in any minimum cotree is equal to $|E| - |V| + 1$, and it is called the *cycle rank* (sometimes the *Betti number*) of G and denoted $\beta(G)$.

By an *adjacency matching* in a subgraph of G, we mean a matching such that each edge in the subgraph is paired with an adjacent edge. For example, one maximum adjacency matching in the cotree of Figure 3.1 contains pairs (a, c) and (b, d), with cotree edges e and f being unpaired.

The following reorganization of Xuong's methods and rederivation of his results is needed for our construction of a maximum-genus algorithm.

LEMMA 3.1. *If a graph G has a completely-paired minimum cotree, then G has a one-face imbedding.*

Proof. By induction on k, the number of edge pairs in the minimum cotree.

BASE CASE: $k = 0$. In this case the graph G is a tree, and every imbedding has exactly one face.

INDUCTIVE CASE: $k > 0$. As an induction hypothesis, assume that a graph with $k - 1$ pairs of edges in a minimum cotree has a one-face imbedding. We now argue that we can add a new pair of adjacent edges $e = (v, w)$ and $f = (w, x)$ to the one-face imbedded graph in the following manner. First insert edge e into the one face in any way between vertices v and w, thereby splitting the single face in two. Note that vertex w now has corners on both faces. Then insert edge f between some corner of x and a corner of w that lies on a different face (see Figure 3.2). This merges the two faces, thereby resulting in a one-face imbedding of $G + e + f$. $\square$


Figure 3.2. Adding adjacent edges e and f to a one-face imbedding.

LEMMA 3.2. *If a graph G has a minimum cotree with k unpaired edges, then G has an imbedding with at most $k + 1$ faces.*

8

Proof. Obtain a one-face imbedding of the spanning tree edges and paired cotree edges of G by the construction in Lemma 3.1. Add each of the k unpaired edges to that imbedding, creating at most one new face for each edge. $\square$

Lemmas 3.1 and 3.2 are constructive, and given a maximum adjacency matching for a minimum cotree, any reasonable implementation of the construction will run in polynomial-time. A naive upper bound on the running time for a graph with e edges is $O(e^2)$.

LEMMA 3.3. *If a graph G has a one-face imbedding, then it has a completely-paired minimum cotree.*

Proof. By induction on the number of edges, k, in G.

BASE CASE I: $k = |V| - 1$. In this case, the graph G is a spanning tree for itself, the cotree is empty, and trivially all edges are paired.

BASE CASE II: $k = |V|$. In this case, the graph G is a spanning tree plus one extra edge. A spanning tree can only be imbedded with one face, and the addition of the extra edge to such a one-face imbedding must break the face in two. Thus, the graph G can only be imbedded with two faces, and the lemma holds vacuously.

INDUCTIVE CASE I: $k > |V|$, and G has a vertex v of degree one. Consider the graph G' obtained by deleting v and its incident edge $e = (v, w)$ from G. Since G has a one-face imbedding, we can readily construct a one-face imbedding of G' by starting with the one-face imbedding for G and deleting e and v. By induction hypothesis, the graph G' has a minimum cotree C with all its edges paired. Since the edge e must be in any spanning tree of G, C is a completely-paired minimum cotree of G.

INDUCTIVE CASE II: $k > |V|$, and G has no vertex of degree one. Consider the boundary-walk around the single face. There must be an edge $r = (u, v)$ whose two appearances in the walk occur as closely together as the two appearances of any other edge. Give the two appearances of r the labels $\overrightarrow{r}$ and $\overleftarrow{r}$, so as to minimize the length of subwalk α from $\overrightarrow{r}$ to $\overleftarrow{r}$. Subwalk α must contain of at least one edge other than r, or else G would have a vertex of degree one, a contradiction. Similarly, if $\overrightarrow{s}$ is the edge following $\overrightarrow{r}$, then α can not also contain $\overleftarrow{s}$, since the two appearances of edge s would then be closer together than those of edge r. Therefore, the boundary-walk around G's face must be of the form $u\,\overrightarrow{r}\,v\,\overrightarrow{s}\,w\alpha_1 v\,\overleftarrow{r}\,u\alpha_2 w\,\overleftarrow{s}\,v\alpha_3 u$, where $s = (v, w)$ is an edge adjacent to r in G, and α_1, α_2, and α_3 are subwalks. See Figure 3.3.

Delete edges r and s from G to obtain the graph G'. Vertices u and v are connected in G' by edges that appeared in subwalk α_3, and vertices v and w are connected by edges that appeared in

9

Figure 3.3. Boundary-walk of G before deleting edges r and s.

subwalk α_1. Every other vertex in G' appeared in α_1, α_2 or α_3 and is thus connected to u, v, or w by edges in G'. Since those three vertices are all connected, it follows that G' is connected.

By the induction hypothesis G' has a cotree C that is completely paired. Clearly the tree $G' - C$ is also a spanning tree of G. Edges r and s can be paired and added to C to form a completely-paired minimum cotree of G. $\square$

LEMMA 3.4.   *If a graph G has a $(k + 1)$-face imbedding, then it has a minimum cotree with at most k unpaired edges in its maximum adjacency matching.*

Proof. By induction on the number k.

BASE CASE: $k = 0$. This follows from the previous lemma.

INDUCTIVE CASE: $k > 1$. Let e be an edge in G that lies on two different faces in some $(k + 1)$-face imbedding. The graph $G - e$ is connected, for otherwise e would lie on only one face, and it has a k-face imbedding when edge e is deleted from the $(k + 1)$-face imbedding of G. By the induction hypothesis, the graph $G - e$ has a minimum cotree C with at most $k - 1$ unpaired edges. Thus $C + e$ is a minimum cotree of G with at most k unpaired edges. $\square$

A *Xuong cotree* of graph G is any minimum cotree of G that admits an adjacency matching with number of **paired edges** maximized (over all minimum cotrees). The number of unpaired edges in such a cotree is **denoted** $U(G)$.

Although Xuong seemed to be little concerned with algorithms, Theorem 3.5 is essentially contained in [X79a]. Theorem 3.6, which relates maximum genus to deficiency, is generally regarded as Xuong's main result.

THEOREM 3.5. [X79A]   *A connected graph G has maximum genus*

$$\gamma_M(G) = \frac{\beta(G) - U(G)}{2}.$$

10

Furthermore, given a Xuong cotree C and a maximum adjacency matching of C, an imbedding of G that minimizes facecount (and thereby maximizes genus) can be found in polynomial-time.

Proof. Follow the construction in Lemma 3.2 to obtain, from a maximum adjacency matching of a Xuong cotree of G, an imbedding with $U(G) + 1$ faces. Lemma 3.4 shows that such an imbedding minimizes the number of faces. Therefore, this is a maximum-genus imbedding in which, by Euler's polyhedral equation, $\gamma_M(G) = (\beta(G) - U(G))/2$. $\square$

THEOREM 3.6. [X79A] *Let G be a connected graph. The maximum genus of G is given by the formula*

$$\gamma_M(G) = \frac{\beta(G) - \xi(G)}{2}.$$

Proof. Any Xuong cotree must contain at least as many odd components as $\xi(G)$. Since every odd component must contain at least one unpaired edge, it follows that $U(G) \geq \xi(G)$. Conversely, a maximum adjacency matching in the complement C of a Xuong tree must leave at least $U(G)$ edges unpaired. The even components of C are completely pairable, and the odd components are pairable except for one edge left over, therefore it follows that $\xi(G) \geq U(G)$. Thus $\xi(G) = U(G)$, from which Xuong's equation follows. Moreover, we see that Xuong trees and Xuong cotrees, as defined here, are indeed complementary objects. $\square$

4. Reduction of Maximum Genus to Linear Matroid Parity

In order to determine the maximum genus and find a maximum imbedding for an arbitrary graph G in polynomial-time, we have shown that it suffices to show how to find a Xuong cotree and a maximum adjacency matching of its edges in polynomial time. This problem resembles what is known as the matroid parity problem for cographic matroids. We use the definitions relating to matroid parity that are found in Stallman and Gabow's paper on linear matroid parity [SG].

A *matroid* $M = (E, I)$ consists of a finite ground set E and a family I of "independent" subsets of E satisfying the following axioms:

1. If $A \in I$ and $B \subseteq A$, then $B \in I$.
2. If $A, B \in I$ and $|A| = |B| + 1$, then there exists $a \in A$ such that $B + a \in I$.

The *matroid parity problem* [La71] is the following. Given a matroid $M = (E, I)$ and a perfect pairing of the elements of the ground set E, find an optimum subset of E such that an element is in the subset if and only if its paired edge is in the subset. Optimum means either a largest subset (the *cardinality parity problem*) or a maximum weighted subset (the *weighted parity problem*). Both can be solved in polynomial time for a large class of matroids known as *linear (or matric)* matroids

11

[Lo,SG,OGHV,S]. The most efficient known algorithm for the cardinality parity problem on general linear matroids runs in $O(nm^3)$ time, where $m = |E|$ and n is the size of the optimum subset. Matroid parity is a generalization of two well-known problems: graph matching and matroid intersection.

For any graph $G = (V, E)$, there is a linear matroid $M = (E, I)$, called the *cographic matroid*, in which the ground set is the edge set of the graph and $C \subseteq E$ is an independent set if and only if $G - C$ is connected. Maximum independent sets in cographic matroids are minimum cotrees of the corresponding graph. For any perfect matching of the edges of the graph, we have an instance of the matroid parity problem on cographic matroids, which we call the *cotree parity problem*. The cardinality parity problems for both graphic (spanning tree) and cographic matroids are easier than general linear matroid parity, and can be solved in $O(nm^2)$ time [La76,SG]. Stallman and Gabow conjecture that this time bound can be reduced to $O(mn \log n)$.

If each edge of a graph G were adjacent to exactly one other edge, then we could directly apply an algorithm for cotree parity to graph G. However, adjacency is not an unambiguous pairing rule for most graphs. Therefore, in this section, we shall transform G into an auxiliary graph G' with unambiguous pairs. The auxiliary graph G' is a subdivision of the graph G itself. Precisely, each edge of G is subdivided into as many edges as its number of edge-neighbors in G. Figure 4.1 illustrates such a subdivision.


Figure 4.1. A graph G and a corresponding auxiliary graph G'.

As illustrated in Figure 4.1, we label each edge of the subdivided graph G' by a label of the form xy, where x is the name of the edge in G of which it is a segment and where y is the name of some distinct neighbor of edge x in the original graph G. The choice of which segment of G is to be labeled xy, for any particular adjacent edge y, is completely arbitrary, provided there is exactly

12

one segment of x labeled xy.

We now consider edge xy to be paired with edge yx. Since this matching is unambiguous, we can apply a cotree parity algorithm to G' and construct—in polynomial time—a minimum cotree C' with a maximum number of paired edges.

Let T' be the edge-complement of the cotree C' in the auxiliary graph G'. Since T' is a spanning tree for the auxiliary graph G', it contains either all the segments or all but one of the segments of every edge of the original graph G. We now associate with spanning tree T' in graph G' a subgraph T in G, according to the rule that an edge x of G appears in T if and only if every segment of x in G' occurs in T'. It is a consequence of the construction of G', T' and T that T is a spanning tree for G: T is acyclic and connected because T' is acyclic and connected.

Let the edge-complement of spanning tree T in the original graph G be called C. Then C is a minimum cotree. Two edges of C are matched if and only if they have matched segments in the cotree C' of the auxiliary graph G'.

This adjacency matching of the edges of cotree C in G is a maximum matching among all possible minimum cotrees of G, because there is a bijection from minimum cotrees of G to minimum cotrees of G' such that the size of the maximum adjacency matching in the cotree of G equals the size of the maximum labeled-edge pairing in G'.

Thus, we have constructed a Xuong cotree for G and a maximum adjacency pairing of its edges in polynomial time.

5. The Algorithm

We now summarize and analyze the algorithm for obtaining a maximum-genus imbedding. Suppose graph G has v vertices, e edges, and maximum degree d. The following steps are used.

1. Create auxiliary graph G' by subdividing edges in G. The new graph has $e' = O(ed)$ edges and $v' = O(ed)$ vertices. This step runs in time $O(ed)$.

2. Run the cotree parity algorithm on graph G'. Producing a maximum set of paired cotree edges requires at most $O(e'(v')^2) = O(e^3 d^3)$ time by the Stallman-Gabow algorithm [SG]. The edge set can be extended to a full minimum cotree by greedily adding unpaired edges. This requires $O((e')^2) = O(e^2 d^2)$ time.

3. For each cotree edge in G', label the corresponding edge in G as a cotree edge. Pair the edges in G which correspond to paired edges in G'. This requires $O(e') = O(ed)$ time.

4. Find a one-face imbedding of the spanning tree edges of G. This requires $O(e)$ time.

5. Add the paired cotree edges to the imbedding. The first edge of each pair can be added in

constant time, but $O(e)$ time is required to find the two resulting faces and determine the placement of the second edge relative to the first. This step requires a total of $O(e^2)$ time.

6. Add unpaired cotree edges to the imbedding. This takes $O(v)$ time, since there is at most one unpaired edge per vertex.

The entire algorithm takes $O(e^3 d^3)$ time. We use the cotree parity algorithm on a special class of graphs containing many vertices of degree two, hence the actual time complexity of this algorithm may be lower. Furthermore, Stallman and Gabow [SG] conjecture that the actual time complexity of their algorithm for general cotree parity is $O(ev \log v)$, which would imply a $O(e^2 d^2 \log ed)$ time bound on our maximum-genus algorithm.

6. Open Problems

1. The fact that maximum genus is reducible to linear matroid parity, which is a generalization of maximum matching, suggests that the corresponding counting problem may be provably difficult. Is it possible that counting the number of ways a graph may be imbedded in a surface of maximum genus is #P-complete?

2. Our algorithm for computing a maximum genus imbedding runs in time polynomial in the size of the graph. This is the only algorithm we know of for constructing any kind of imbedding that runs in time independent of the genus. Is it possible to extend the algorithm to return imbeddings in which the genus is a fixed constant less than the maximum?

3. Reif [R79] showed that determining whether a partial imbedding of a graph can be extended to a full imbedding in the same surface is NP-complete. Is the same true for determining whether a partial imbedding can be extended to a one-face imbedding?

4. Suppose graphs G and H are non-isomorphic. One might ask how the non-isomorphism shows up in the way the graphs may be imbedded in different surfaces. Knowing all the "counting information" about how a graph imbeds in all surfaces is not a complete invariant for isomorphism. It clearly isn't a complete invariant for trees, which only have planar imbeddings, and we have examples of non-isomorphic, highly connected, graphs such that counting the number of imbeddings in all surfaces does not distinguish them. However, randomly sampling imbeddings and making estimates of the number of ways different graphs imbed in different surfaces may prove to be an interesting new isomorphism heuristic.

7. References

[E] J. Edmonds, "A combinatorial representation for polyhedral surfaces." *Notices Amer. Math. Soc.*, 7, 1960, p. 646.

[F] I. Filotti, "An algorithm for imbedding cubic graphs in the torus." *J. Comp. Sys.*, 20, 1980, pp. 255-276.

[FMR] I. Filotti, G. Miller, J. Reif, "On determining the genus of a graph in $O(v^{O(G)})$ steps." 11th *Annual ACM Symposium on Theory of Computing*, 1979, pp. 27-37.

[GF] J. Gross, M. Furst, "Hierarchy for imbedding-distribution invariants of a graph." Typescript, 1985.

[GR] J. Gross, R. Rosen, "A linear-time planarity algorithm for 2-complexes." *Journal of the ACM*, 20, 1979, pp. 611-617.

[GT] J. Gross, T. Tucker, *Topological Graph Theory*. Wiley-Interscience, to appear.

[H] L. Heffter, "Über das Problem der Nachbargebiete." *Math. Ann.*, 38, 1891, pp. 477-508.

[HT] J. Hopcroft, R. Tarjan, "Efficient planarity testing." *Journal of the ACM*, 21, 1974, pp. 549-568.

[J] M. Jungerman, "A characterization of upper embeddable graphs." *Trans. Amer. Math. Soc.*, 241, 1978, pp. 401-406.

[La71] E. Lawler, "Matroids with parity conditions: a new class of combinatorial optimization problems." Electronics Research Laboratory, Berkeley. Memo. ERL-M334, 1971.

[La76] E. Lawler, *Combinatorial Optimization, Networks, and Matroids*. Holt, Rinehart, and Winston, 1976, 374 pages.

[Lo] L. Lovász, "The matroid matching problem." *Algebraic Methods in Graph Theory*, for the Colloquia Mathematica Societatis Janos Bolyai, Szeged, Hungary, 1978.

[N] L. Nebeský, "Every connected, locally connected graph is upper embeddable." *J. Graph Theory*, 5, 1981, pp. 197-199.

[NSW] E. Nordhaus, B. Stewart, A. White, "On the maximum genus of a graph." *Journal of Combinatorial Theory*, Series B, 11, 1971, pp. 258-267.

[OGHV] J. Orlin, E. Gugenheim, J. Hammond, J. VandeVate, "Linear matroid parity made almost easy." Extended abstract.

[R79] J. Reif, "The complexity of extending a graph imbedding." Unpublished manuscript, 1979.

[R70] R. Ringeisen, *The Maximum Genus of a Graph.* Ph. D. Thesis, Michigan State University, 1970.

[R72] R. Ringeisen, "$K_{m,n}$ has 2-cell imbeddings." *Journal of Combinatorial Theory,* Series B, 12, 1972, pp. 101-104.

[S] M. Stallman, "Weighted matroid parity." Unpublished manuscript, 1984.

[SG] M. Stallman, H. Gabow, "An augmenting path algorithm for the parity problem on linear matroids." 25th *Symposium on the Foundations of Computer Science,* 1984. pp. 217-227.

[W] A. White, *Graphs, Groups, and Surfaces* (2nd ed.). North-Holland/Elsevier Science, 1984, 314 pages.

[X79a] N. Xuong, "How to determine the maximum genus of a graph." *Journal of Combinatorial Theory*, Series B, 26, 1979, pp. 216-225.

[X79b] N. Xuong, "Upper-embeddable graphs and related topics." *Journal of Combinatorial Theory*, Series B, 26, 1979, pp. 226-232.

[Z] J. Zaks, "The maximum genus of cartesian products of graphs." *Canad. J. Math.*, 26, 1974, pp. 1025-1035.