

Generating Maps on Surfaces

Thom Sulanke¹ 

Received: 1 September 2015 / Revised: 15 December 2016 / Accepted: 19 December 2016 /
Published online: 13 January 2017
© Springer Science+Business Media New York 2017

Abstract We describe procedures for generating all 2-cell embedded simple graphs with up to a fixed number of vertices on a given surface. We also modify these procedures to generate closed 2-cell embeddings and polyhedral embeddings. We give results of computer implementations of these procedures for seven surfaces: the sphere, the torus, the double torus, the projective plane, the Klein bottle, the triple cross surface, and the quadruple cross surface.

Keywords Map · Irreducible triangulation · Open 2-cell embedding · Closed 2-cell embedding · Polyhedral embedding

Mathematics Subject Classification 05C10 · 05C30 · 05C85

1 Introduction

A *map* is a simple graph embedded in a surface such that every face is simply connected. A *triangulation* is a map in which every face has three edges. Section 2 contains more detailed definitions. In Sects. 3, 4, and 5 we describe the operations or local deformations which we apply to triangulations and maps.

The generation of triangulations and maps on surfaces such as the projective plane, the torus, and the Klein bottle has similarities to the generation of triangulations and maps on the sphere. There are also interesting differences.

Editor in Charge: Günter M. Ziegler

Thom Sulanke
tsulanke@indiana.edu

¹ Department of Physics, Indiana University, Bloomington, IN 47405, USA




Fig. 1 Expansion for splitting a vertex




Fig. 2 Expansion for removing an edge

Brinkmann and McKay provide procedures for generating triangulations and maps on the sphere [4,5]. They implement these procedures in a computer program *plantri* [3]. We extend these techniques to non-spherical surfaces.

We know from the work of Steinitz [11] that we can generate all the triangulations of the sphere with $n > 4$ vertices by applying the vertex splitting operation (Fig. 1 and Sect. 5) to the triangulations of the sphere with $n - 1$ vertices. The single initial triangulation for this recursive process is K_4 embedded in the sphere, the (boundary of the) tetrahedron. The inverse of the vertex splitting operation is the operation of edge contraction. To assure that the repeated application of the vertex splitting operation generates all triangulations it is required that the edge contraction operation can always be performed on any triangulation other than K_4 and that the result of this edge contraction operation is also a triangulation. For a triangulation of the sphere other than K_4 it is indeed always possible to find an edge for which the edge contraction operation can be applied [11].

Triangulations of non-spherical surfaces can be generated in the same way. However, for any surface other than the sphere there are many but a finite number of triangulations which do not have any contractible edges [2]. An edge is not contractible if an attempt to apply the edge contraction operation would produce multiple edges. We define a triangulation with no contractible edges to be an *irreducible triangulation*. The initial triangulations used for generating all triangulations of a surface are the irreducible triangulations of this surface. In Sect. 7 we provide a method for producing the class of irreducible triangulations of a surface.

We next turn to the generation of maps which might not be triangulations. Maps on the sphere can be obtained from the triangulations of the sphere by the operation of edge removal (Fig. 2 and Sect. 4). For the repeated application of the operation of edge removal to generate all the maps on the sphere, it must always be possible to perform the inverse operation of adding an edge to a map which is not a triangulation and, in the process, obtain another map. The edge must be added in such a way that does not create multiple edges. By applying the Jordan curve theorem we can show that for the sphere the operation of adding an edge is always possible for a map which is not a triangulation.

We would like to generate maps on other surfaces in a similar way. However, there are maps which are not triangulations for which it is not possible to add an edge without producing multiple edges. In Sect. 5 we introduce a class of maps, which we call *irreducible maps*, which are analogous to irreducible triangulations for the purpose of generating maps.

The generation of the maps with n vertices of a fixed surface consists of four steps:

1. Generate the *irreducible triangulations* of the surface (Sect. 7).
2. Generate the *irreducible maps* of the surface from the irreducible triangulations by removing vertices (Sect. 6).
3. Split vertices (E-expansions) of the irreducible maps to obtain *face irreducible maps* with n vertices (Fig. 1 and Sect. 5).
4. Remove edges (D-expansions) of the face irreducible maps while the maps remain *2-cell embeddings* (Fig. 2 and Sect. 4).

Before describing these steps in Sects. 4–7 we give definitions related to graphs in Sect. 2 and we provide notation used when describing operations for generating maps in Sect. 3. The computer program *surftri* [14] is an implementation of the techniques for generating triangulations and maps described in this paper. Information on how this was done is contained in Sect. 8. In Sect. 9 we discuss how the steps for generating maps can be modified to generate closed 2-cell embeddings and polyhedral embeddings. We display the irreducible maps on the projective plane and torus in Sect. 10.

2 Definitions

A *surface* is a two-dimensional compact manifold. We denote the orientable surface with genus g , the sphere with g handles attached, as S_g and the nonorientable surface with genus g , the sphere with g crosscaps attached, as N_g .

We consider only simple graphs which are graphs with no loops and no multiple edges. Let G be a connected graph embedded on a surface S . A *face* of G is a connected component of the complement of G in S . The graph G is a *map* (or *open 2-cell embedding*) on S if every face of G is homeomorphic to an open disk. If the three edges v_1v_2 , v_2v_3 , and v_3v_1 are contained in the map G then the union of these three edges is a *3-cycle* of G denoted as $v_1v_2v_3$. A map G on a surface S is a *triangulation* of S if the boundary of every face of G is a 3-cycle and the map is not a single 3-cycle embedded on the sphere.

Let F be a face of G with the boundary edges, in order, v_1v_2 , v_2v_3 , ..., v_mv_1 . The face F is denoted by the list of vertices $v_1v_2v_3 \dots v_m$. These vertices do not need to be distinct. We call F an *m-face* and we say that F has *size* m . If F has size $m \geq 4$ then F is defined to be a *large face*. The subscripts of the vertices of a face are modulo m . We say we *triangulate* a large face F when we add vertices and edges to the interior of F in such a way that all the new faces formed are 3-faces.

3 Recursive Generation

The basic “isomorph-free” generation technique that is used for the Steps 1–4 is described in detail in [4, 5, 8]. For each generation process we specify the class \mathcal{C}

which is the class of maps to be generated, the initial class $\mathcal{C}_0 \subseteq \mathcal{C}$ which is the class of maps from which the maps in \mathcal{C} are generated, and F an expansion operation. The expansion operation F is a function from \mathcal{C} into the set of subclasses of \mathcal{C} . We say $(\mathcal{C}_0; F)$ generates \mathcal{C} if for each $G \in \mathcal{C}$ there is a sequence $G_0, G_1, \dots, G_m = G$ such that $G_0 \in \mathcal{C}_0$ and for every i , $1 \leq i \leq m$, we have $G_i \in F(G_{i-1})$. We call the expansion operation F the F -expansion. The inverse of the expansion operation F is the F -reduction.

Figures 1 and 2 represent expansion operations which we use to generate maps. The left side of each figure shows a part of the embedded graph before the expansion operation. The right side shows the same part of the graph after the expansion operation. Each part of the graphs shown is contained in a simply connected component of the surface. The full edges which are shown are required to be a part of the map being modified. The shorter half edges (in Fig. 1) are other unchanged edges of the map. The small flattened triangles (in Figs. 1, 2) represent the location of zero or more other unchanged edges. The expansion operation replaces the subgraph on the left with the subgraph on the right. The reduction operation replaces the right subgraph with the left subgraph. Since the part of the graph being modified is contained in a simply connected component the surface remains the same.

4 Removing Edges in Corners of Large Faces

We examine the steps of the overall map generation process in reverse order to help clarify the choice of operations and initial classes which we use.

In Step 4 we generate \mathcal{M}_2 , the class of all maps for a surface S . We consider an arbitrary map of the type being generated and describe the reduction operation, the D-reduction. The D-expansion is the inverse of the D-reduction. Let $\mathcal{M}_1 \subseteq \mathcal{M}_2$ be the class of those maps on S for which the D-reduction is not possible. We characterize \mathcal{M}_1 and show that $(\mathcal{M}_1; D)$ generates \mathcal{M}_2 .

Let G be a map on a surface S . We continue to call the map G even as it is modified by the D-reduction. The D-expansion which is shown in Fig. 2 is the removal of an edge from a 3-face. It can only be applied if the map has a 3-face. The D-reduction is the addition of an edge in the “corner” of a large face. We only apply the D-reduction if the edge to be added does not already exist in the map.

Let $F = v_1 v_2 \dots v_m$ be a large face of G and let v_i be a vertex on F . If $v_i = v_{i+2}$ then v_{i+1} has degree 1 and is not adjacent to v_{i+3} . The edge $v_{i+1} v_{i+3}$ can be added in the interior of F dividing F into a 3-face and a face of size $m-1$. Adding this edge reduces the number of vertices of degree 1. Repeated addition of edges of this type results in no faces having a vertex of degree 1.

We can now assume that the three vertices v_i, v_{i+1}, v_{i+2} on F are distinct. If v_i is not adjacent to v_{i+2} then again the edge $v_i v_{i+2}$ can be added in the interior of F dividing F into a 3-face and a face of size $m-1$.

The addition of an edge in this way is the D-reduction. Even though a D-reduction adds an edge to the map it simplifies the map by making it more like a triangulation by increasing the number of 3-faces. The application of a D-reduction reduces the size

of one large face and does not change the size of the other large faces. So eventually no more D-reductions are possible.

If S is the sphere then a D-reduction can always be applied in a large face. The initial class \mathcal{M}_1 of maps on the sphere for which the D-reduction is not possible is the class of triangulations of the sphere.

If S is a surface other than the sphere then it is possible to have a large face in which it is not possible to use a D-reduction. For example, P^4 , in Fig. 5 shows K_4 embedded in the projective plane. Opposite points on the hexagon have been identified. The map has three 4-faces. None of the faces are 3-faces and no edges can be added since the graph is complete.

A large face $F = v_1 v_2 \dots v_m$ of a map is an *irreducible face* if for every i , $1 \leq i \leq m$, the vertices v_i and v_{i+2} are adjacent. A map is *face irreducible* if every face is an irreducible face or a 3-face. Trivially, every triangulation is a face irreducible map.

So the initial class \mathcal{M}_1 of maps on S for which the D-reduction is not possible is the class of face irreducible maps on S .

5 Splitting Vertices

We now consider Step 3 of the map generation process in which we generate the face irreducible maps. The class of maps to be generated is \mathcal{M}_1 which is used as the initial class in the previous section. The E-expansion shown in Fig. 1 is used. Below we specify an initial class \mathcal{M}_0 such that $(\mathcal{M}_0; E)$ generates \mathcal{M}_1 .

Let $G \in \mathcal{M}_1$ be a face irreducible map on the surface S . The E-reduction, *edge contraction*, is the inverse of the E-expansion. The E-reduction is applied only if the faces on both sides of the edge being contracted are 3-faces. Also the E-reduction is performed only if the resulting graph is still simple. Let v_1 and v_2 be the vertices of the edge to be contracted and let $u_1 v_1 v_2$ and $u_2 v_1 v_2$ be the 3-faces on either side of $v_1 v_2$. The two ends of the contracted edge, v_1 and v_2 , must not both be adjacent to any vertices other than u_1 and u_2 . Otherwise, multiple edges would be produced when $v_1 v_2$ is contracted. An edge is *contractible* if it is on exactly two 3-cycles both of which are 3-faces. To apply the edge contraction operation to an edge the edge must be contractible and the map must not be K_4 embedded in the sphere. An edge is *noncontractible* if there is at least one 3-cycle containing the edge which is not a 3-face. An edge is *essentially noncontractible* if at least one 3-cycle containing the edge is an essential 3-cycle on S .

Theorem 5.1 *Every edge on the boundary of an irreducible face is essentially noncontractible.*

Proof Let $F = v_1 v_2 \dots v_m$ be an irreducible face and let $v_i v_{i+1}$ be an edge on the boundary of F . Define $C_{F,i}$ to be the 3-cycle $v_i v_{i+1} v_{i+2}$. Assume $C_{F,i}$ is not essential, i.e. the interior of $C_{F,i}$ is simply connected. Let D be the disk consisting of $C_{F,i}$ and its interior. The path $v_i v_{i+1} v_{i+2}$ is on the boundary of both F and D . The graph consisting of the vertices and edges in D is a map on the sphere so F cannot be in D . The vertex v_{i+3} must be adjacent to v_{i+1} but the edge $v_{i+3} v_{i+1}$ cannot be in the interior of F or in the interior of D . So $v_{i+3} = v_i$. Similarly, $v_{i+2} = v_{i-1}$. Thus the edge

$e = v_{i+2}v_{i+3} = v_{i-1}v_i = v_{i+2}v_i$ must occur twice on F and F is on both sides of e . This is impossible since $C_{F,i}$ must be on one side of e . \square

Contracting edges of G does not change the surface S in which G is embedded so $C_{F,i}$ in the proof of Theorem 5.1 remains essential and cannot become a 3-face when edges are contracted. So an irreducible face of G remains an irreducible face of G with the same size when edges of G are contracted. As edges of G are contracted the number of irreducible faces and their sizes remain unchanged.

We apply the E-reduction while contractible edges remain in G . Each application of the E-reduction reduces the number of vertices of G . So after a finite number of E-reductions there are no contractible edges.

We define a map G as an *irreducible map* if G is face irreducible and no edge in G is contractible. The following property of irreducible maps is a generalization of a similar property of irreducible triangulations.

Theorem 5.2 *Every edge of an irreducible map is essentially noncontractible.*

Proof Let G be an irreducible map. An edge of G is either on the boundary of an irreducible face or on two 3-faces. In the first case the edge is essentially noncontractible by Theorem 5.1. In the second case since the edge is not contractible it must be on a 3-cycle C which is not a 3-face. We show that C is essential. Assume C is not essential. On the surface C bounds a disk D . A new map H on the sphere can be obtained by replacing the exterior of D with a 3-face.

The map H cannot contain a large face. If H contains a large face $F = v_1v_2 \dots v_m$ then at least one vertex, say v_2 , of F must be in the interior of the (assumed) non-essential 3-cycle C . Then the essential 3-cycle $v_1v_2v_3$ is in the disk D . But this is not possible.

The map H contains no large face and is a triangulation of the sphere. If H is K_4 then any interior edge of D is contractible in G . If H is not K_4 then there are contractible edges of H which are also contractible in G . \square

The initial class \mathcal{M}_0 is the class of all irreducible maps on the surface S and $(\mathcal{M}_0; E)$ generates \mathcal{M}_1 , the class of all face irreducible maps on the surface S .


6 Generating Irreducible Maps

Irreducible maps which we generate in Step 2 have a very nice property which makes them easy to obtain. Each irreducible map G is a “submap” of an irreducible triangulation T , i.e. G is obtained by removing zero or more vertices and adjacent edges from T .

The class to be generated is the class of irreducible maps. The initial class is the class of irreducible triangulations. The expansion operation is to remove a set of vertices and adjacent edges with the condition that the resulting embedded graph is in the class of irreducible maps. Theorem 6.2 shows that this expansion operation does generate all of the irreducible maps.

The following theorem proves another property of irreducible faces and provides us with the tool to prove Theorem 6.2.

Fig. 3 Expansion for removing a vertex, $k = 4$



Theorem 6.1 *Let F be an irreducible face of a map G (which need not be irreducible). We may replace the interior of F with new vertices, edges, and faces to create a new map G' such that the new faces are 3-faces, the new edges are noncontractible in G' , and the interior of the union of the new vertices, edges, and faces is simply connected.*

Proof There is one special irreducible face which we handle separately. P^3 shown in Fig. 5 is a single 3-cycle embedded in the projective plane. The dotted lines and open circles in the figure represent the edges and vertices which are added to obtain the irreducible triangulation P^1 . So we now can assume that F is not the one face of P^3 .

Let $F = v_1v_2 \dots v_m$ be an irreducible face of G . We could triangulate F by placing an m -cycle $w_1w_2 \dots w_m$ in the interior of F and an additional vertex x inside this m -cycle. Adding the edges w_iv_i, w_iv_{i+1}, w_ix for $i, 1 \leq i \leq m$, would fill F with 3-faces. We could then contract edges in the interior of F until no more edges in the interior of F are contractible. However, this map might not contain G . This might occur if, while contracting edges in the interior of F , two vertices of the boundary of F are the ends of an edge which is contracted thus merging these two vertices of G . So triangulating of F must be done with more care.

We use the V-reduction to add one vertex at a time to the interior of F and to attach the new vertex to vertices which are on the boundary of the large face. The V-reduction is shown in Fig. 3. This figure is similar to the previous figures showing expansions and reductions. Figure 3 shows only four vertices of the large face, more vertices may be used. The figure also shows dashed curves. Each dashed curve represents part of an edge which is not completely contained in the simply connected component of the surface represented by the figure. One requirement for the application of the V-reduction is that at least four edges are added joining consecutive vertices on the boundary of the face to the new vertex. With each V-reduction the size of the large face is reduced. We can apply the V-reduction only a finite number of times. Let the new vertex be w and let these consecutive vertices on the boundary of the face be v_1, v_2, \dots, v_k with $k \geq 4$. A second requirement for the application of the V-reduction is that the 3-cycles $v_iv_{i+1}v_{i+2}$ be essential for $i, 1 \leq i \leq k - 2$. Since the 3-cycle wv_iv_{i+2} is homeomorphic in S to the essential 3-cycle $v_{i+1}v_iv_{i+2}$, for $i, 1 \leq i \leq k - 2$ each new edge wv_i is essentially noncontractible for $i, 1 \leq i \leq k$.

We now describe how to add the first vertex to the interior of F .

First assume the vertices v_1, v_2, \dots, v_m of F are distinct. This is always the case when $m \leq 5$ since each vertex on F is adjacent to all the other vertices on F . Since F is an irreducible face and $m \geq 4$ we use all the vertices of F for the V-reduction.

The added edges are all essentially noncontractible and all the new faces are 3-faces satisfying the theorem.

Now assume the vertices of F are not distinct. There is at least one vertex which occurs at least twice on the boundary of F . To triangulate the interior of F we must add more than one new vertex to the interior of F to prevent multiple edges.

Let k be the number of vertices in the longest paths on the boundary of F . Since there is at least one vertex which occurs at least twice on the boundary of F we have $k < m$. If necessary we relabel the vertices so that $p = v_1 v_2 \dots v_k$ is one of the longest paths on the boundary of F . The vertices v_1, v_2, \dots, v_k are all distinct. The vertex v_m occurs at least twice on F and is on p for, otherwise, we could extend p to v_m producing a longer path on the boundary of F . Since v_m is adjacent to v_1 and v_2 , $v_m = v_a$ for some a , $3 \leq a \leq k$. Likewise, $v_{k+1} = v_b$ for some b , $1 \leq b \leq k - 2$.

We assert that $k \geq 4$. Since F is an irreducible face v_i, v_{i+1} , and v_{i+2} are pairwise adjacent and distinct for every i , $1 \leq i \leq m$. So $k \geq 3$. Suppose that $k = 3$. Then for every i , $1 \leq i \leq m$, we have $v_i = v_{i+3}$. The edges of F would be $v_1 v_2, v_2 v_3, v_3 v_1, v_1 v_2, v_2 v_3, v_3 v_1, \dots$. Since a face can occur on an edge at most twice $F = v_1 v_2 v_3 v_1 v_2 v_3$ and the map is P_3 contrary to our earlier assumption.

We use $\{v_1, v_2, \dots, v_k\}$ to apply a V-reduction and call the new vertex w_1 . The resulting large face F_1 is $v_1 w_1 v_k \dots v_m$. We observe that the face F_1 might not be irreducible but it almost satisfies the definition. For every i , $k \leq i \leq m - 1$, v_i is adjacent to v_{i+2} . Since $v_m = v_a$ and $v_{k+1} = v_b$, w_1 is adjacent to v_m and v_{k+1} . The only possible missing condition is that v_1 might not be adjacent to v_k .

In the remainder of the proof we add additional vertices w_2, w_3, \dots , in a similar way. With each additional vertex w_{n+1} a face F_{n+1} is produced which has fewer vertices of the original face F than F_n has. When a face F_{n+1} with no vertices of the original face F is obtained then we finish triangulating F as described below.

Assume we have added n new vertices to the interior of F and we have a large face F_n with boundary $v_1 w_1 \dots w_n v_j \dots v_m$ such that $j \leq m$ and w_n is adjacent to v_{j+1} (which is v_1 if $j = m$). We have shown above that this assumption is true for $n = 1$. Starting from this assumption for n we show either (i) that we can finish triangulating F or (ii) that the assumption is true for $n + 1$ and the face F_{n+1} has fewer vertices of the original face F than F_n has. Thus the construction terminates in a finite number of steps.

We obtain (i) when the vertices on the boundary of F_n are distinct. In this case, the vertices for attaching w_{n+1} are $\{w_n, v_j, \dots, v_m, v_1, w_1\}$. Recall that when w_1 is attached w_1 is adjacent to v_m . If $n + 1 \leq 3$ then there is no longer a large face in F and all the edges which have been added in the interior of F are essentially noncontractible. If $n + 1 > 3$ then there is a resulting large face F_{n+1} with boundary $w_1 w_2 \dots w_{n+1}$. We arbitrarily triangulate the face F_{n+1} with edges $w_{n+1} w_i$ for i , $1 < i < n$. Some of these $n - 2$ edges might be contractible but all the other edges which have been added to the interior of F are essentially noncontractible. We repeatedly contract any contractible edge in the interior of F_{n+1} until there are no contractible edges in the interior of F_{n+1} . All of the edges in the interior of F are then noncontractible and we have triangulated the face F as required.

We can show (ii) when the vertices on the boundary of F_n are not distinct. The vertices w_n, v_j, v_{j+1} , and v_{j+2} are distinct since w_n is not on the boundary of F

and the vertices v_j , v_{j+1} , and v_{j+2} are pairwise adjacent. Let $v_{j'}$ be the vertex on the boundary of F_n such that $w_n, v_j, v_{j+1}, v_{j+2}, \dots, v_{j'}$ are distinct and $v_{j'+1}$ is in $\{v_j, v_{j+1}, v_{j+2}, \dots, v_{j'-2}\}$. We use $\{w_n, v_j, v_{j+1}, v_{j+2}, \dots, v_{j'}\}$ to apply a V-reduction with the new vertex w_{n+1} . In this way we obtain a smaller face F_{n+1} which fulfills our assumption for $n + 1$ which is (ii). \square

Theorem 6.2 *If G is an irreducible map on a surface S then there is at least one irreducible triangulation T of S from which G may be obtained by removing a set of vertices from T along with the edges containing these vertices.*

Proof Let G be the irreducible map. Using Theorem 6.1 we “irreducibly triangulate” each irreducible face of G to produce a triangulation T . The edges of T which are in G are essentially noncontractible by Theorem 5.2. The edges of T which are not in G are noncontractible by Theorem 6.1. It may be possible to triangulate the irreducible faces of G in more than one way so the irreducible triangulation T may be one of many which satisfy the theorem. \square

7 Generating Irreducible Triangulations

Irreducible triangulations which are generated in Step 1 have been extensively studied. For any fixed surface the number of irreducible triangulations is finite [2]. Irreducible triangulations have been determined and displayed by a number of authors: the single irreducible triangulation of the sphere (S_0) by Steinitz and Rademacher [11]; the two irreducible triangulations of the projective plane or the cross surface (N_1) by Barnette [1]; the 21 irreducible triangulations of the torus (S_1) by Lawrencenko [6]; and the 29 irreducible triangulations of the Klein bottle (N_2) by Lawrencenko and Negami [7] and Sulanke [13]. The irreducible triangulations of the double torus (S_2), the triple cross surface (N_3), and the quadruple cross surface (N_4) have been generated by the author using an extension of computer program *surftri* [14]. The counts of irreducible triangulations are shown in Table 1. The largest of these classes (N_4) required 54 CPU days. The author estimates it would take CPU centuries for this program to generate the irreducible triangulations for S_3 or N_5 .

We describe briefly how the two stage generation process for irreducible triangulations of a surface S works. In the first stage we use the vertex splitting operation to generate triangulations on slightly simpler surfaces than S . We impose certain conditions necessary for the second stage on these triangulations. These conditions limit the triangulations generated to a finite number. In the second stage these triangulations are modified in such a way that new handles or crosscaps are added to produce irreducible triangulations on S . Only the reduction operation for the second stage is described. More details may be found in [12].

Let S (not the sphere) be the surface for which we are generating irreducible triangulations. Let G be an irreducible triangulation of S . Theorem 7.1 below shows that G contains many nonseparating 3-cycles. Let $w_1 w_2 w_3$ be a nonseparating 3-cycle of G . We create a new triangulation G' of a different surface S' using the operations described below.

Table 1 The number of irreducible triangulations and irreducible maps

	Vertices	S_1	S_2	N_1	N_2	N_3	N_4	
Irreducible triangulations	6			1				
	7	1		1				
	8	4			6			
	9	15			19	133	37	
	10	1	865		2	2521	10,347	
	11		26,276		2	4638	370,170	
	12		117,047			1320	1,891,557	
	13		159,205			946	2,067,817	
	14		54,527			93	956,967	
	15		38,195			50	700,733	
	16		664			7	186,999	
	17		5				89,036	
	18						19,427	
	19						3975	
	20						832	
	21						79	
	22						6	
	Irreducible triangulations	Total	21	396,784	2	29	9708	6,297,982
	Irreducible maps with large faces	3			1			
		4	2		1	3	3	
		5	7	34	2	14	111	441
		6	11	885	1	28	886	15,059
7		15	7522		43	3531	146,463	
8		11	36,395		38	9285	784,407	
9		1	124,890		14	17,229	2,872,513	
10			314,021		4	18,703	7,473,401	
11			526,436			10,124	13,025,842	
12			480,596			4678	13,565,448	
13			208,029			1026	9,468,143	
14			82,073			282	5,557,358	
15			3382			30	2,292,654	
16			24				862,799	
17							222,153	
18						47,833		
19						7978		
20						636		
21						30		
Irreducible maps with large faces	Total	47	1,784,287	5	144	65,888	56,343,158	
Irreducible maps	Total	68	2,181,071	7	173	75,596	62,641,140	

If the 3-cycle $w_1w_2w_3$ in G is two-sided then we cut along $w_1w_2w_3$ to produce a surface S' with a boundary consisting of two disjoint 3-cycles $u'_1u'_2u'_3$ and $v'_1v'_2v'_3$ where u'_i and v'_i come from the original vertex w_i for $i = 1, 2, 3$. We cap the holes with two 3-faces $u'_1u'_2u'_3$ and $v'_1v'_2v'_3$. G' is now a triangulation of S' which has an Euler genus two less than the Euler genus of S .

If the 3-cycle $w_1w_2w_3$ in G is one-sided then we cut along $w_1w_2w_3$ to produce a surface S' with a boundary consisting of the 6-cycle $u'_1u'_2u'_3v'_1v'_2v'_3$ where u'_i and v'_i again come from the original vertex w_i for $i = 1, 2, 3$. We cap the hole with a new vertex t' and six 3-faces $t'u'_1u'_2$, $t'u'_2u'_3$, $t'u'_3v'_1$, $t'v'_1v'_2$, $t'v'_2v'_3$, and $t'v'_3u'_1$. G' is now a triangulation of S' which has an Euler genus one less than the Euler genus of S .

The following theorem is similar to Lemma 4 of [1] and Lemma 4 of [6]. In a triangulation the *link* of a vertex v is the cycle which is the boundary of the union of the faces containing v .

Theorem 7.1 *Let G be an irreducible triangulation of a surface other than the sphere, let v be a vertex of G , and let L be the link of v . Then there are two nonseparating 3-cycles vv_iv_k and vv_jv_l such that v_i, v_j, v_k , and v_l are distinct and one path from v_i to v_k in L contains v_j and the other path from v_i to v_k in L contains v_l .*

Proof Since G is irreducible, for any vertex u in L the edge vu is on a nonfacial 3-cycle vuw . Pick two vertices v_i and v_k in L for which vv_iv_k is a nonfacial 3-cycle and the distance from v_i to v_k in L is minimal. The shorter path from v_i to v_k in L must have an interior vertex since vv_iv_k is not a face. Let the vertex v_j be such an interior vertex on the shorter path from v_i to v_k in L . Let v_l be a vertex in L such that vv_jv_l is a nonfacial 3-cycle. v_l is not on the path from v_i to v_k in L containing v_j since the distance from v_j and v_l in L is at least the distance from v_i and v_k in L . Suppose vv_iv_k separates the surface. Then v_j and v_l would be in different components but v_jv_l is an edge. Therefore, vv_iv_k is nonseparating and, similarly, vv_jv_l is also nonseparating. \square

8 Implementation

The computer program *surftri* [14] implements the procedures to generate maps on various surfaces. Many of the ideas and much of the code used in *surftri* are taken from the work of Brinkmann and McKay. Their program *plantri* [3] generates triangulations and maps on the sphere as well as other classes of planar graphs.

The operation of *plantri* is described in [4, 5]. To obtain triangulations of the sphere with n vertices *plantri* starts with the only irreducible triangulation of the sphere, K_4 , and vertices are split using variations of the E-expansion until the triangulations have n vertices. If maps with n vertices are being generated then as each triangulation with n vertices is produced the program switches to the mode of using the operation of removing edges. Edges are removed from 3-faces using the D-expansion.

To generate triangulations of a surface in the program *surftri* we start with the irreducible triangulations of that surface. Vertices are split using the E-expansion to obtain triangulations of the surface using procedures similar those used in *plantri*. The list of the irreducible triangulations is provided as input to *surftri*.

To generate maps on a surface with n vertices *surftri* generates the face irreducible maps with n vertices. Vertex splitting, the E-expansion, is again used. *surftri* starts with irreducible maps and vertex splitting is only done if each irreducible face remains irreducible. The irreducible maps are provided as input, having been pregenerated and stored on disk. As each face irreducible map with n vertices is produced *surftri* switches to the mode of using the operation of removing edges. Edges are removed from 3-faces using the D-expansion.

The irreducible maps with at least one large face were generated using Theorem 6.2 rather than the construction used in its proof. For a fixed surface each irreducible triangulation was processed by removing sets of vertices and checking if the results were irreducible maps. Duplicates were removed by sorting all the irreducible maps obtained in this way. The number of irreducible triangulations of a surface is finite [2] and each irreducible map is obtained by removing vertices from an irreducible triangulation. Therefore, the number of irreducible maps on a surface is finite. We also show the counts of irreducible maps in Table 1.

Not only can we implement the procedures described in this paper in computer software but we can also use the generated initial classes and the expansion operations to provide inductive proofs.

As a simple example, we determine the minimum number of vertices in a map on the double torus. By examining the irreducible maps on the double torus we see that each one has at least five vertices. The E-expansion operation increases the number of vertices by one while the D-expansion operation leaves the number of vertices unchanged. Thus a map on the double torus has at least 5 vertices.

9 Maps, Closed 2-Cell Embeddings, and Polyhedral Embeddings

We have described the steps listed in Sect. 1 for generating maps (open 2-cell embeddings). We now consider how more restricted classes of maps can be generated by modifying these steps.

The *face-width* of an embedded graph on a surface is the smallest number k such that there is a noncontractible closed curve on the surface that intersects the graph at k points [10].

Maps are embedded graphs which have face-width at least 1 and are 1-connected.

A *closed 2-cell embedding* is a map for which the closure of every face is a closed 2-cell. For every face F of a closed 2-cell embedding no vertex occurs more than once on F . Closed 2-cell embeddings are those maps which have face-width at least 2 and are 2-connected [9]. The D-expansion does not increase the face-width or the connectivity of a map. So in order to obtain a closed 2-cell embedding when we apply the D-expansion we must apply the operation to another closed 2-cell embedding. However, the E-expansion may increase the face-width of some maps from 1 to 2. We only modify the final step used to generate maps to obtain a procedure for generating closed 2-cell embeddings.

The generation of the closed 2-cell embeddings with n vertices of a fixed surface consists of four steps:

Table 2 Counts by number of vertices for maps on the projective plane

Vertices	Irreducible maps		Face irreducible maps		Maps				
	Tri.	Closed 2-cell	Open 2-cell	Triangular	Closed 2-cell	Open 2-cell	Polyhedral	Closed 2-cell	Open 2-cell
3	0	0	1	0	0	1	0	0	1
4	0	1	1	0	1	2	0	1	6
5	0	2	2	0	2	4	0	9	65
6	1	2	2	1	5	11	1	188	1128
7	1	1	1	3	18	42	10	4850	27,041
8	0	0	0	16	90	204	247	141,255	741,730
9	0	0	0	134	566	1185	8576	4,138,394	21,858,099
10	0	0	0	1210	4059	7768	263,539	119,621,784	663,735,117
11	0	0	0	11,719	31,773	55,832	7,290,012	3,413,905,527	20,513,612,018
12	0	0	0	114,478	261,912	427,171	185,392,686	96,571,642,059	640,173,121,863
13	0	0	0	1,108,826	2,222,281	3,398,145	4,448,447,624	2,716,506,356,624	–
14	0	0	0	10,606,795	19,187,942	27,792,593	102,469,338,545	–	–
15	0	0	0	100,352,404	167,528,480	231,869,302	2,292,650,424,419	–	–
16	0	0	0	940,956,644	1,473,754,923	1,963,615,254	–	–	–
17	0	0	0	8,762,227,629	13,035,546,948	1,682,269,5958	–	–	–
18	0	0	0	81,168,427,279	115,780,306,285	145,453,671,164	–	–	–
19	0	0	0	748,953,936,818	1,031,742,846,949	1,267,008,314,524	–	–	–

Table 3 Counts by number of vertices for maps on the torus

Vertices	Irreducible maps			Face irreducible maps			Maps		
	Tri.	Closed 2-cell	Open 2-cell	Triangular	Closed 2-cell	Open 2-cell	Polyhedral	Closed 2-cell	Open 2-cell
	4	0	0	2	0	0	2	0	0
5	0	2	7	0	2	8	0	3	70
6	0	7	11	0	8	25	0	205	2656
7	1	12	16	1	39	124	1	15,958	126,466
8	4	15	15	7	290	789	33	1,014,018	6,070,817
9	15	16	16	112	2584	5976	4713	52,587,939	280,232,378
10	1	1	1	2109	25,202	49,677	442,429	2,376,996,732	12,389,481,487
11	0	0	0	37,867	258,518	443,088	28,635,972	97,845,502,685	527,699,182,180
12	0	0	0	605,496	2,719,039	4,145,672	1,417,423,218	3,770,598,166,962	—
13	0	0	0	8,778,329	28,922,902	40,158,242	58,321,972,887	—	—
14	0	0	0	117,839,254	308,507,103	398,230,859	2,102,831,216,406	—	—
15	0	0	0	1,491,505,713	3,283,573,624	4,008,448,713	—	—	—
16	0	0	0	18,035,839,188	34,773,632,099	40,687,327,452	—	—	—
17	0	0	0	210,391,127,053	365,879,925,813	414,537,434,014	—	—	—




Fig. 4 Irreducible triangulations of the Projective Plane




Fig. 5 Irreducible maps with large faces on the Projective Plane

1. Generate the *irreducible triangulations* of the surface.
2. Generate the *irreducible maps* of the surface from the irreducible triangulations by removing vertices.
3. Split vertices (E-expansions) of the irreducible maps to obtain *face irreducible maps* with n vertices.
4. Remove edges (D-expansions) of the face irreducible maps while the maps remain *closed 2-cell embeddings*.

A *polyhedral embedding* is a map for which the closures of any pair of faces have exactly one vertex, exactly one edge, or no points in common. Polyhedral embeddings on a surface are those maps which have face-width at least 3 and are 3-connected [9].

We use Theorem 5.1 to show that any map with an irreducible face has face-width at most 2. Let $F = v_1 v_2 \dots v_m$ be an irreducible face on a surface S . We construct a closed curve consisting of two segments. One segment is in the interior of F connecting




Fig. 6 Irreducible triangulations of the torus, T^1 – T^{12}

v_1 and v_3 . The other segment is close to the edge v_1v_3 and connects them. This closed curve is homeomorphic in S to the essential 3-cycle $v_1v_2v_3$ and thus is noncontractible. This closed curve intersects the graph at only 2 points, v_1 and v_3 .




Fig. 7 Irreducible triangulations of the torus, T^{13} – T^{21}

Again we note that the D-expansion does not increase the face-width nor the connectivity of a map of a map. Also, the E-expansion is restricted from eliminating irreducible faces. So we do not need to use any maps with irreducible faces in the generation of polyhedral embeddings.

The generation of the polyhedral embeddings with n vertices of a fixed surface consists of only three steps:

1. Generate the *irreducible triangulations* of the surface.
2. Split vertices (E-expansions) of the irreducible triangulations to obtain *triangulations* with n vertices.
3. Remove edges (D-expansion) of the triangulations while the maps remain *polyhedral embeddings*.

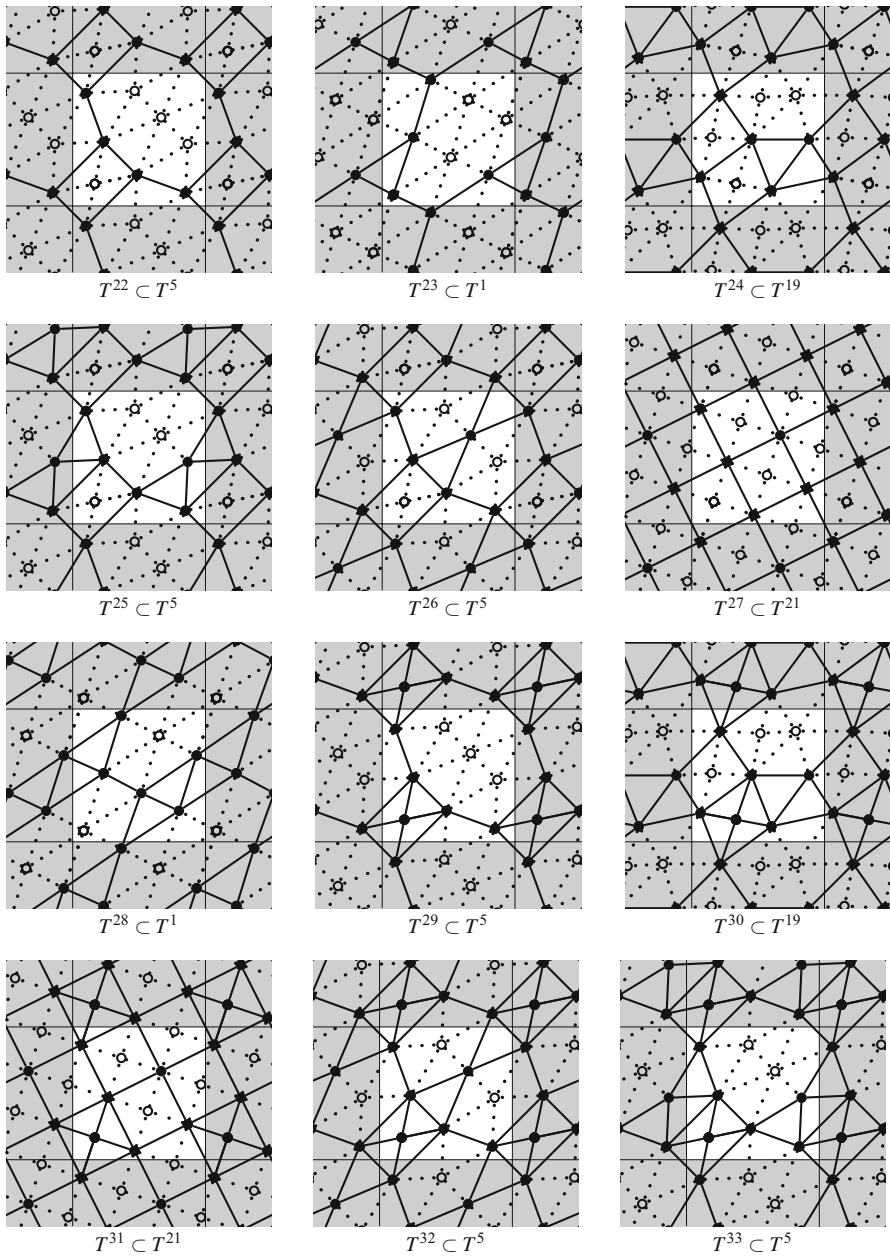


Fig. 8 Irreducible maps with large faces on the torus, T^{22} – T^{33}

In Table 2 we show the counts of maps, closed 2-cell embeddings, and polyhedral embeddings on the projective plane for increasing numbers of vertices. We also show the counts for irreducible maps and face irreducible maps. The *surfri* program produced these values. Table 3 shows the counts for maps on the torus.

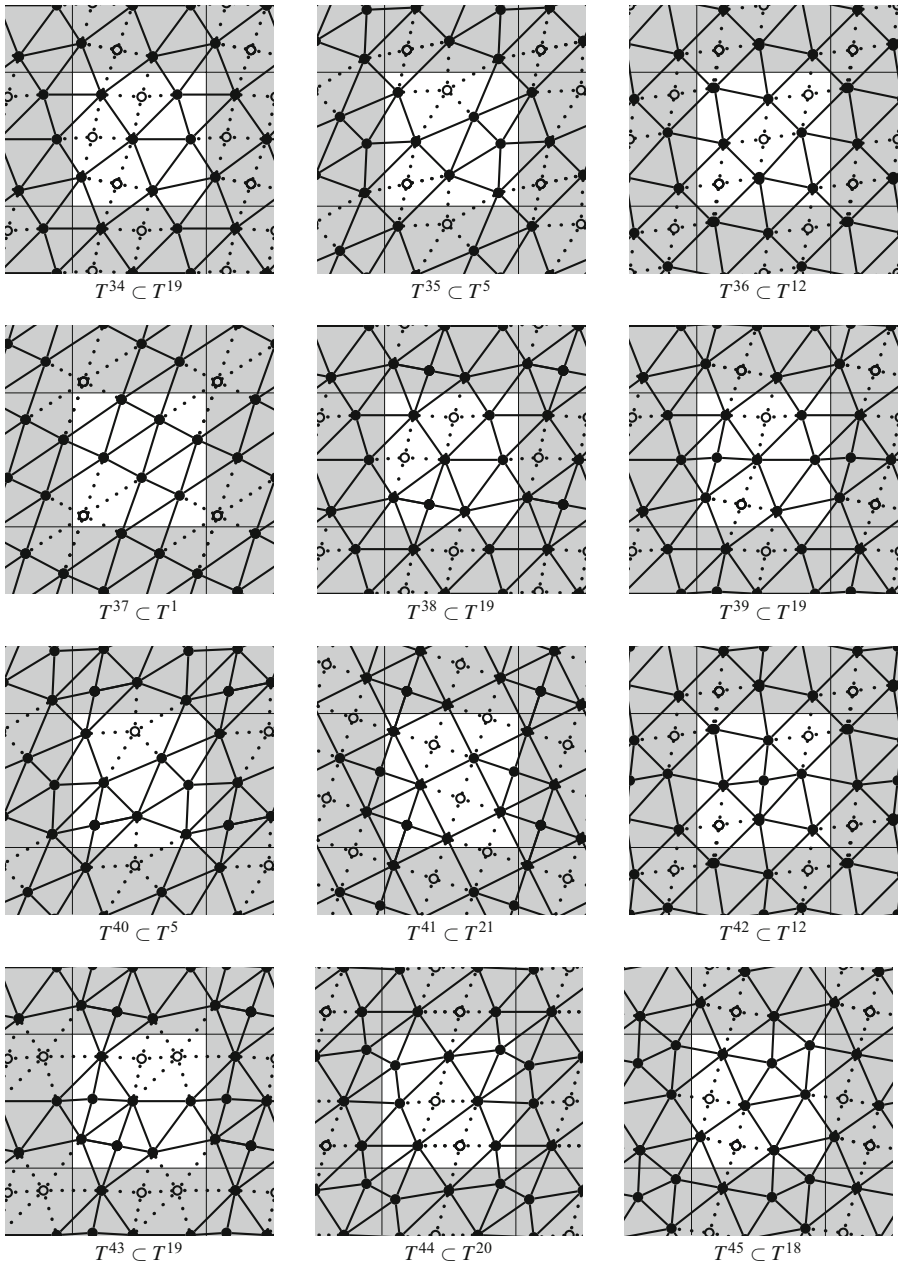


Fig. 9 Irreducible maps with large faces on the torus, T^{34} – T^{45}




Fig. 10 Irreducible maps with large faces on the torus, $T^{46} \text{--} T^{57}$

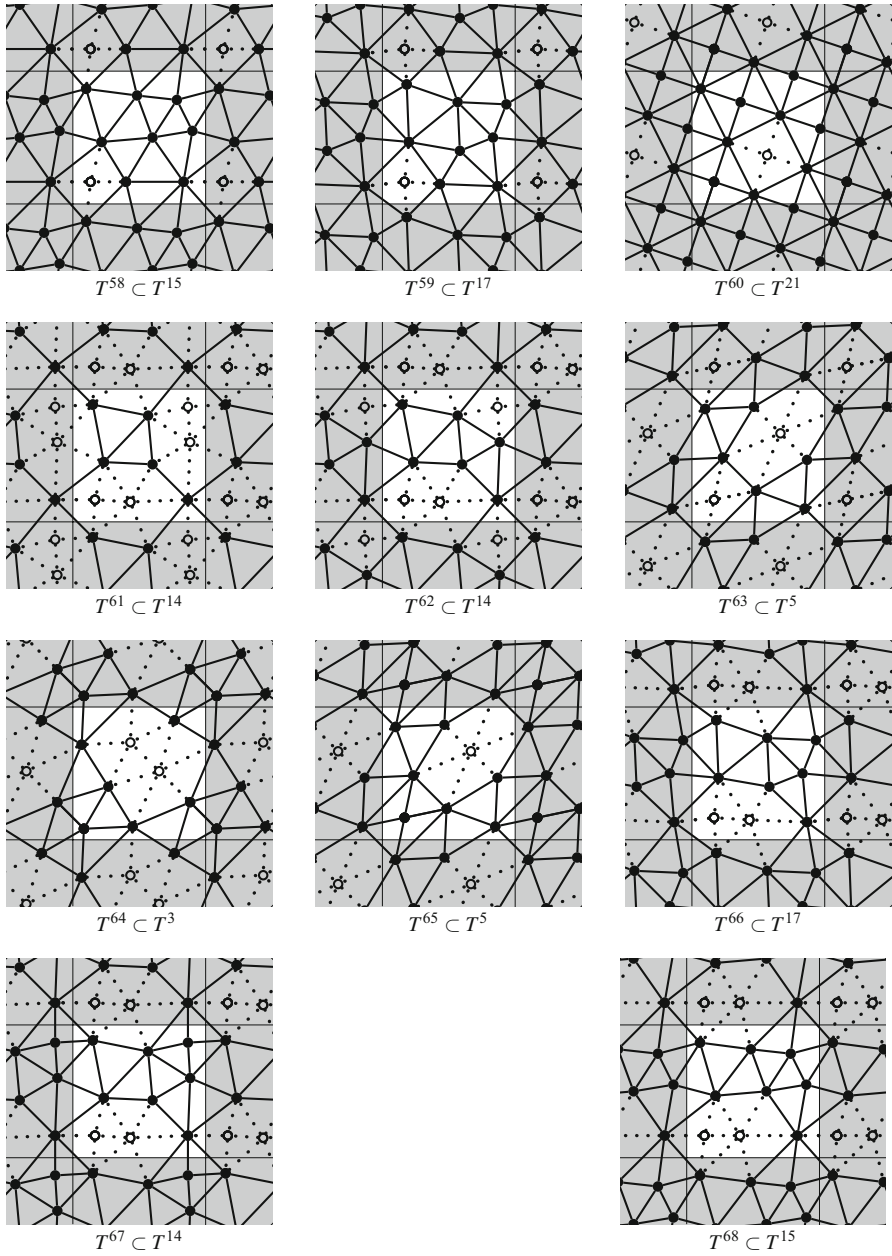


Fig. 11 Irreducible maps with large faces on the torus, T^{58} – T^{68}

We estimate that the entries in the tables which are blank would require more than 100 days of CPU time to compute on 2.4 GHz processors. We adapted the data structures used in the *surftri* program to store the embeddings from the data structures used in *plantri*. We modified these data structures to allow embeddings in non-orientable surfaces. These modified data structures require more computer operations than are used in *plantri*. The generation rates for *surftri* range from 1 to 1.4 million maps/second on a 2.4 GHz processor. When generating maps on the sphere the rates for *surftri* are 0.7–0.95 of those rates observed using *plantri*.

10 Irreducible Maps on the Projective Plane and the Torus

We provide drawings of the irreducible maps on the projective plane and torus. For each irreducible map with large faces we indicate one of the irreducible triangulations from which it may be obtained. The vertices and edges which have been removed from the irreducible triangulation are shown as open circles and dotted lines. The irreducible triangulation chosen requires the minimum number of vertices to be removed.

Figure 4 shows the two irreducible triangulations of the projective plane [1]. Figure 5 shows the five irreducible maps with large faces on the projective plane.

There are 21 irreducible triangulations, T^1 – T^{21} , of the torus [6] which are shown in Figs. 6 and 7. Figures 8, 9, 10 and 11 show the 47 irreducible maps, T^{22} – T^{68} , with large faces on the torus.

References

1. Barnette, D.: Generating the triangulations of the projective plane. *J. Comb. Theory Ser. B* **33**(3), 222–230 (1982)
2. Barnette, D.W., Edelson, A.L.: All 2-manifolds have finitely many minimal triangulations. *Israel J. Math.* **67**(1), 123–128 (1989)
3. Brinkmann, G., McKay, B.D.: Plantri, program for generation of certain types of planar graph. <http://cs.anu.edu.au/~bdm/plantri/>
4. Brinkmann, G., McKay, B.D.: Fast generation of planar graphs. *MATCH Commun. Math. Comput. Chem.* **58**(2), 323–357 (2007)
5. Brinkmann, G., McKay, B.D.: Fast generation of planar graphs (expanded version) (2008). <http://cs.anu.edu.au/~bdm/papers/plantri-full.pdf>
6. Lavrenchenko, S.A.: Irreducible triangulations of a torus. *Ukrain. Geom. Sb.* **30**, 52–62 (1987)
7. Lawrencenko, S., Negami, S.: Irreducible triangulations of the Klein bottle. *J. Comb. Theory Ser. B* **70**(2), 265–291 (1997)
8. McKay, B.D.: Isomorph-free exhaustive generation. *J. Algorithms* **26**(2), 306–324 (1998)
9. Mohar, B., Thomassen, C.: *Graphs on Surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore (2001)
10. Robertson, N., Seymour, P.D.: Graph minors. VII. Disjoint paths on a surface. *J. Comb. Theory Ser. B* **45**(2), 212–254 (1988)
11. Steinitz, E., Rademacher, H.: *Vorlesungen über die Theorie der Polyeder unter Einschluss der Elemente der Topologie*. Grundlehren der Mathematischen Wissenschaften, vol. 41. Springer, Berlin (1976)
12. Sulanke, T.: Generating irreducible triangulations of surfaces. [arXiv:math.CO/0606687](https://arxiv.org/abs/math/0606687)
13. Sulanke, T.: Note on the irreducible triangulations of the Klein bottle. *J. Comb. Theory Ser. B* **96**(6), 964–972 (2006)
14. Sulanke, T.: Source for surftri and lists of irreducible triangulations. <http://hep.physics.indiana.edu/~tsulanke/graphs/surftri/>