



HAL
open science

Irreducible Triangulations of Surfaces with Boundary

Alexandre Boulch, Éric Colin de Verdière, Atsuhiko Nakamoto

► **To cite this version:**

Alexandre Boulch, Éric Colin de Verdière, Atsuhiko Nakamoto. Irreducible Triangulations of Surfaces with Boundary. *Graphs and Combinatorics*, 2013, 29 (6), pp.1675-1688. 10.1007/s00373-012-1244-1 . hal-01163747

HAL Id: hal-01163747

<https://hal.science/hal-01163747>

Submitted on 16 Jun 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Irreducible Triangulations of Surfaces with Boundary*

Alexandre Boulch[†]

Éric Colin de Verdière[‡]

Atsuhiko Nakamoto[§]

Abstract

A triangulation of a surface is *irreducible* if no edge can be contracted to produce a triangulation of the same surface. In this paper, we investigate irreducible triangulations of surfaces with boundary. We prove that the number of vertices of an irreducible triangulation of a (possibly non-orientable) surface of genus $g \geq 0$ with $b \geq 0$ boundary components is $O(g + b)$. So far, the result was known only for surfaces without boundary ($b = 0$). While our technique yields a worse constant in the $O(\cdot)$ notation, the present proof is elementary, and simpler than the previous ones in the case of surfaces without boundary.

1 Introduction

Let \mathcal{S} be a surface, possibly with boundary. A *triangulation* is a simplicial complex whose underlying space is \mathcal{S} . Contracting an edge of the triangulation (identifying two adjacent vertices in the simplicial complex) is allowed if this results in another triangulation of the same surface. A triangulation is *irreducible* (or *minimal*) if no edge is contractible. Every triangulation can be reduced to an irreducible triangulation by iteratively contracting some of its edges.

Irreducible triangulations have been much studied in the context of surfaces without boundary. In this paper, we extend known results to the case of surfaces with boundary. Specifically, we prove that the number of vertices of an irreducible triangulation is linear in the genus and the number of boundary components of the surface. Compared to previous works, our theorem and its proof have two interesting features: the result is more general, since it applies to surfaces with boundary, and the arguments of the proof are simpler.

1.1 Previous Works for Surfaces Without Boundary

We first describe previous related works, on surfaces without boundary. Barnette and Edelson [5, 6] proved that the number of irreducible triangulations of a given surface is finite. Nakamoto and Ota [26] were the first to show that the number of vertices in an irreducible triangulation is at most linear in the genus of the surface. The best upper bound known to date is due to Joret and Wood [15], who proved that this number is at most $\max\{13g - 4, 4\}$. (Here and in the sequel, g is the *Euler genus*, which equals twice the usual genus for orientable surfaces and equals the usual genus for non-orientable surfaces.) This bound is asymptotically tight, as there are irreducible triangulations with $\Omega(g)$ vertices; however, the minimal number of vertices in a triangulation is $\Theta(\sqrt{g})$ [16].

Some low genus cases were studied. Steinitz [31] proved that the unique irreducible triangulation of the sphere is the boundary of the tetrahedron. The two irreducible triangulations of the projective plane were found by Barnette [4], followed by the 21 irreducible triangulations of the torus by Lawrencenko [19] and the 29 triangulations of the Klein bottle by Sulanke [35]. More recently, Sulanke [33, 34] developed a method to generate

*This work was done during the first author's internship at École normale supérieure. The internship was funded by the *Agence Nationale de la Recherche* under the *Triangles* project of the *Programme blanc* ANR-07-BLAN-0319.

[†]École Polytechnique (member of ParisTech), Palaiseau, France, alexandre.boulch@polytechnique.edu

[‡]Laboratoire d'informatique, École normale supérieure, CNRS, Paris, France, eric.colin.de.verdiere@ens.fr

[§]Department of Mathematics, Yokohama National University, Yokohama, Japan, nakamoto@ynu.ac.jp

all the irreducible triangulations of surfaces without boundary. His algorithm rediscovered the irreducible triangulations for the projective plane, the Klein bottle, and the torus; it also built the irreducible triangulations of the double-torus (396,784 triangulations) and the non-orientable surfaces of genus three (9,708) and four (6,297,982).

Generalizations of the notion of irreducible triangulations, such as *k-irreducible triangulations*, also called *k-minimal triangulations* ($k \geq 3$), have also been studied [13, 20]. Juvan et al. [17, Section 6] also study this concept in the case of surfaces with boundary; their proof technique implies that the number of irreducible triangulations is finite for every surface (possibly with boundary). For a more detailed survey on results on irreducible triangulations, see Mohar and Thomassen [22, Sect. 5.4]. Higher-dimensional analogs have also been studied, and in particular conditions ensuring that contracting an edge of a simplicial complex preserves the topological type [3, 10].

1.2 Applications of irreducible triangulations

One motivation for studying irreducible triangulations is that, to solve some problems on triangulations, it sometimes suffices to solve them on irreducible triangulations. For example, on a triangulation of an orientable surface with Euler genus $g \geq 4$ (at least two handles), Barnette [22, Conjecture 5.9.3] conjectured that there always exists a cycle without repeated vertices that is non-null-homotopic and separating. More generally, Mohar and Thomassen [22, Conjecture 5.9.5] conjectured that for every even h , $0 < h < g$, there exists a cycle without repeated vertices that splits the surface into two surfaces of genus h and $g - h$, respectively. To prove these conjectures, it suffices to prove them for irreducible triangulations. (See also the discussion by Sulanke [34, Sect. 5].)

Irreducible triangulations have also been used to characterize triangulations without K_6 -minors. (The characterization of abstract graphs without K_m -minors has been done for any $m \leq 5$, but the problem for $m \geq 6$ seems to be very difficult.) The key fact for the characterization is that every triangulation on a surface \mathcal{S} with no K_6 -minor is transformed into an irreducible triangulation with no K_6 -minor by contracting edges. The complete lists of irreducible triangulations are known only for surfaces of Euler genus at most four, and so the characterizations are done only for those surfaces [18, 23–25].

Irreducible triangulations can be used to generate all triangulations of a given surface [29]. They are also a good tool to study diagonal flips on triangulations. Negami [27, 28] used the fact that there are finitely many irreducible triangulations to prove that two triangulations of a surface with the same number of vertices are equivalent under diagonal flips, provided the number of vertices is greater than an integer depending only on the surface. For further applications, see the recent paper by Joret and Wood [15] and references therein.

1.3 Our Result

It turns out that the notion of irreducible triangulations extends directly to the case of surfaces with boundary. In this paper, we prove that the number of vertices of such an irreducible triangulation admits an upper bound that is linear in the genus g and the number of boundaries b of the surface. In more detail:

Theorem 1. *Let \mathcal{S} be a (possibly non-orientable) surface with Euler genus g and b boundaries. Assume $g \geq 1$ or $b \geq 2$. Then every irreducible triangulation of \mathcal{S} has at most $570g + 385b - 573$ vertices, except for the case of the projective plane ($g = 1$ and $b = 0$), in which the bound is 186.*

This bound is asymptotically tight; see Figure 1. Compared to the case of surfaces without boundary, the main difficulty we encountered was to prove that the number of boundary vertices is $O(g + b)$ (there are indeed irreducible triangulations of surfaces whose single boundary contains $\Theta(g)$ vertices, as Figure 1 also illustrates); the known methods for surfaces without boundary do not seem to extend easily to surfaces with boundary. Our strategy is roughly as follows. Let T be an irreducible triangulation. First, we show that every matching of (the vertex-edge graph of) T has $O(g + b)$ vertices. Then, we show that every inclusionwise maximal matching contains a constant fraction of the vertices of T . For technical reasons, in the case of surfaces with boundary, we actually need to restrict ourselves to a matching satisfying some additional mild conditions.

In particular, we reprove that, on a surface without boundary of genus g , the number of vertices of an irreducible triangulation is $O(g)$. Our method does not improve over the current best bound of $\max\{13g - 4, 4\}$ by Joret and Wood [15]. However, it is substantially different and simpler than the other known proofs of this




Figure 1: For any even $g \geq 0$ and any $b \geq 1$ (one of these two inequalities being strict), there exists an irreducible triangulation of an orientable surface with Euler genus g with b boundary components, and with $5g/2 + 4b - 2$ vertices. The figure illustrates the case $g = 4$ and $b = 3$. Starting with a set of triangles glued together, all meeting at a vertex (bottom part), attach a set of $g/2$ pairs of interlaced rectangular strips (top left) and a set of $b - 1$ non-interlaced rectangular strips (top right), and triangulate every strip by adding an arbitrary diagonal (not shown in the picture). That the resulting triangulation is irreducible follows from the fact that every edge belongs to a non-null-homotopic 3-cycle or is a linking edge (a non-boundary edge whose endpoints are both boundary vertices). Also, note that all vertices are on the boundary. In particular, taking $b = 1$, we obtain an irreducible triangulation whose single boundary component contains $5g/2 + 2$ vertices.

result. These former proofs, by Nakamoto and Ota [26] and Joret and Wood [15] (see also Gao et al. [12]), rely on a deep theorem by Miller [21] (see also Archdeacon [1]) stating that the genus of a graph (the minimum Euler genus of a surface on which a graph can be embedded) is additive over 2-vertex amalgams (identification of two vertices of disjoint graphs). Another paper by Cheng et al. [9] also claims a linear bound without using Miller's theorem, but this part of their paper has a flaw (personal communication with the authors).¹ In contrast, our proof is short and uses only elementary topological lemmas.

Finally, we refine the above technique in the case of surfaces without boundary, and obtain a bound that is better than that of Theorem 1, but no better than the current best result by Joret and Wood [15].

We shall introduce some definitions from topology and preliminary lemmas in Section 2. We then prove our main theorem (Section 3). Finally, in Section 4, we describe the improvement of the technique for surfaces without boundary.

2 Preliminaries

We present a few notions of combinatorial topology; for further details, see also Stillwell [32], Armstrong [2], or Henle [14].

2.1 Topological Background

Surfaces, Cycles, and Homotopy. A *surface* (2-manifold with boundary) is a topological Hausdorff space where each point has an open neighborhood homeomorphic to the plane or the closed half-plane; the points in the latter case are called *boundary points*. Henceforth, \mathcal{S} denotes a compact, connected surface.

By the classification theorem, \mathcal{S} is homeomorphic to a surface obtained from a sphere by removing finitely many open disks and attaching handles (*orientable case*) or Möbius bands (*non-orientable case*) along some of the resulting boundaries. In the orientable case, the *Euler genus* of \mathcal{S} , denoted by g , equals *twice* the number of

¹Specifically, in the proof of their Lemma 3, the authors incorrectly claim that there are at most g pairwise non-homologous cycles on an orientable surface of Euler genus g .




Figure 2: Edge contraction on an embedded graph.

handles; in the non-orientable case, it equals the number of Möbius bands. The number of remaining *boundary components* is denoted by b .

In this paper, a *cycle* on \mathcal{S} is the image of a one-to-one continuous map $S^1 \rightarrow \mathcal{S}$, where S^1 is the standard circle. In particular, we emphasize that cycles are undirected and simple. Two cycles are *homotopic* if one can be deformed continuously to the other; more formally, two cycles C_0 and C_1 are homotopic if there exists a continuous map $h : [0, 1] \times S^1 \rightarrow \mathcal{S}$ such that $h(0, \cdot)$ is one-to-one and has image C_0 , and similarly $h(1, \cdot)$ is one-to-one and has image C_1 . A cycle is *null-homotopic* if and only if it bounds a disk on \mathcal{S} . We emphasize that only homotopy of cycles is considered in this paper; for example, we say that two loops are homotopic if and only if the corresponding cycles are homotopic (without fixing any point of the loops).

A cycle is *two-sided* if cutting along it results in a (possibly disconnected) surface with two boundaries, and *one-sided* otherwise. Equivalently, a cycle is two-sided if it has a neighborhood homeomorphic to an annulus, and one-sided if it has a neighborhood homeomorphic to a Möbius band (which implies that the surface is non-orientable). Two cycles in general position *cross* at point p if they intersect at p and the intersection cannot be removed by an arbitrarily small perturbation of the cycles. Two homotopic cycles in general position cross an even number of times if they are two-sided, and an odd number of times if they are one-sided.

Graph Embeddings, Triangulations, and Edge Contractions. Let G be a graph, possibly with loops and multiple edges. An *embedding* of G on \mathcal{S} is a “crossing-free” drawing of G on \mathcal{S} . More precisely, the vertices of G are mapped to distinct points of \mathcal{S} ; each edge is mapped to a path in \mathcal{S} , meeting the image of the vertex set only at its endpoints, and such that the endpoints of the path agree with the points assigned to the vertices of that edge. Moreover, all the paths must be without intersection or self-intersection except, of course, at common endpoints. We sometimes identify G with its embedding on \mathcal{S} . The *faces* of G are the connected components of the complement of the image of G in \mathcal{S} . A graph embedding G is *cellular* if each of its faces is an open disk. If it is the case, *Euler’s formula* states that $|V| - |E| + |F| = 2 - g - b$, where V , E , and F are the sets of vertices, edges, and faces of G , respectively.

Let e be an edge of a graph G embedded in the interior of \mathcal{S} . Assume that e is not a loop. The *contraction* of e shrinks e to a single vertex; the resulting graph is in the interior of \mathcal{S} . Loops and multiple edges may appear during this process (Figure 2).

A *triangulation* T on \mathcal{S} is a graph without loops or multiple edges embedded on \mathcal{S} such that each face is an open disk with three distinct vertices, and such that two such triangles intersect on a single edge (and its two incident vertices), a single vertex, or not at all. In other words, the vertices, edges, and faces of G form a simplicial complex whose underlying space is \mathcal{S} .

The definition of edge contraction on a triangulation is slightly different from an edge contraction on a graph embedding. Let uv be an edge of T ; contracting edge uv identifies both vertices u and v in the simplicial complex T ; the dimension of some simplices decreases by one. We say that uv is *contractible* if the new simplicial complex is still homeomorphic to \mathcal{S} (Figure 3).

In more detail, assume for now that $e = uv$ is an interior edge, incident with triangles uvx and uvy . Contracting e shrinks e , identifying its two vertices u and v , and identifies the two pairs of parallel edges $\{ux, vx\}$ and $\{uy, vy\}$. The definition is similar if e is a boundary edge, except that it has a single incident triangle uvx . If uv is not a boundary edge but exactly one vertex (say u) is incident to a boundary, then the edge uv is contracted to u , on the boundary. If this operation results in a triangulation of \mathcal{S} , we say that e is *contractible*. In particular, a *linking edge* of T is a non-boundary edge whose both vertices are on the boundary; a linking edge is never contractible.

A triangulation of a surface is *irreducible* if it contains no contractible edge. For example, it is known that




Figure 3: Edge contraction on a triangulation.

the only irreducible triangulation of the sphere is the boundary of a tetrahedron [31]. Using a similar argument, it is not hard to show that the only irreducible triangulation of the disk is a single triangle.

2.2 Preliminary Lemmas

We list here a series of basic facts that will be used in our proof.

Lemma 2. *Assume \mathcal{S} is not the sphere or the disk, and let T be an irreducible triangulation of \mathcal{S} . Then every non-linking edge of T belongs to a non-null-homotopic 3-cycle.*

Proof. This was proved by Barnette and Edelson [5, Lemma 1] for surfaces without boundary: In this case, every edge of T belongs to a 3-cycle that is not the boundary of a triangle; if that 3-cycle is null-homotopic, then it bounds a disk, and an edge inside that disk must be contractible. The argument immediately extends to non-linking edges of surfaces with boundary. (For boundary edges, we need to distinguish whether the boundary has length at least four, in which case the previous argument applies, or exactly three, in which case the result is obvious.) \square

Lemma 3. *The degree of a non-boundary vertex of an irreducible triangulation of \mathcal{S} is at least four.*

Proof. This is a direct consequence of a result by Sulanke [33, Theorem 1]. Specifically, he uses Lemma 2 to show that every vertex of an irreducible triangulation belongs to two non-separating 3-cycles crossing at that vertex. Again, the argument extends to non-boundary vertices of surfaces with boundary. \square

Lemma 4. *Let G be a 1-vertex graph with ℓ loop edges, embedded in the interior of \mathcal{S} . Assume that no face of G is a disk bounded by one or two edges. Then $\ell \leq 3g + 2b - 3$, except if \mathcal{S} is a sphere or a disk (in which cases $\ell = 0$).*

Proof. Barnette and Edelson [6, Corollary 1] prove a similar result; see also Chambers et al. [8, Lemma 2.1]. Here is a sketch of proof. Without loss of generality, we can assume that G is inclusionwise maximal; namely, no edge can be added to G without violating the hypotheses of the lemma. Then it follows from the classification of surfaces that, unless the surface is the sphere, the disk, or the projective plane, every face of the graph is a disk bounded by three edges, or an annulus bounded by a single edge and a single boundary component of \mathcal{S} . A standard double-counting argument combined with Euler's formula concludes. \square

Corollary 5. *Let G be a 1-vertex graph with ℓ loop edges, embedded in the interior of \mathcal{S} . Assume that no loop of G is null-homotopic and that no two loops of G are homotopic. Then $\ell = 0$ if \mathcal{S} is a sphere or a disk, $\ell \leq 1$ if \mathcal{S} is a projective plane, and $\ell \leq 3g + 2b - 3$ otherwise.*

Proof. The hypotheses imply that no face of G is a disk bounded by one or two edges (and thus Lemma 4 concludes), unless that disk is bounded by twice the same edge (in which case \mathcal{S} is the projective plane). \square

Lemma 6. *Let C be a non-null-homotopic 3-cycle in an irreducible triangulation T of \mathcal{S} . No more than nine pairwise edge-disjoint 3-cycles of T are homotopic to C .*




Figure 4: Illustration of the proof of Lemma 6 (in the two-sided case): if there are ten homotopic edge-disjoint 3-cycles, there must be four pairwise disjoint homotopic 3-cycles, so there is at least one contractible edge.

Proof. First case: C is two-sided. This case is a small variation on a lemma by Barnette and Edelson [5, Lemma 9]. Any two distinct 3-cycles homotopic to C must cross an even number of times, hence cannot cross at all; thus two such 3-cycles bound an annulus, possibly “pinched” on a vertex or an edge. So the set of 3-cycles homotopic to a given 3-cycle can be ordered linearly. Assume there are at least ten pairwise edge-disjoint 3-cycles of T homotopic to C ; let us consider ten such consecutive cycles in this ordering, C_1, \dots, C_{10} . See Figure 4.

For every i , the annulus between C_i and C_{i+3} cannot be pinched along a vertex: otherwise, it is easy to see that an edge between C_{i+1} and C_{i+2} would be contractible [5, Lemma 7]. This annulus cannot be pinched along an edge, since the cycles are edge-disjoint. So any two consecutive cycles in the sequence C_1, C_4, C_7, C_{10} bound a non-pinched annulus. Now, similarly, some edge between C_4 and C_7 is contractible [5, Lemma 9], which is a contradiction.

Second case: C is one-sided. In this case, any 3-cycle homotopic to C crosses C exactly once, and must therefore share exactly one vertex with C . Let v be any vertex of C ; we prove below that at most two 3-cycles different from C and homotopic to C pass through v . This proves that there are at most seven 3-cycles homotopic to C (including C itself), which concludes.

So assume that (at least) four 3-cycles (including C) homotopic to C share together a vertex v . These cycles lie in a Möbius band “pinched” at v , and we can order them linearly; let C_1, \dots, C_4 be consecutive cycles in this ordering. As in the first case [5, Lemma 7], an edge between C_2 and C_3 would be contractible, a contradiction. \square

3 Proof of Theorem 1

A *matching* M of a graph G is a set of edges of G such that every vertex of G belongs to at most one edge of M .

3.1 The Size of a Matching

Our first task is to prove that a matching of an irreducible triangulation has size $O(g + b)$.

Proposition 7. *Let T be an irreducible triangulation of \mathcal{S} , where $g \geq 1$ or $b \geq 2$. Let M be a matching of T containing no linking edge. Then the number of edges of M is at most 27 if \mathcal{S} is the projective plane ($g = 1$ and $b = 0$) and $81g + 54b - 81$ otherwise.*

Proof. The structure of the proof is as follows. In three steps, we remove edges from M , obtaining successive matchings M_1, M_2 , and M_3 , each of them satisfying additional properties. We show that the edge set of M_3 is in bijection with the edge set of a 1-vertex graph on \mathcal{S} where no edge is null-homotopic and no two edges are homotopic. By Corollary 5, this implies that M_3 has $O(g + b)$ edges. Furthermore, we show that M_3 contains some constant fraction of the edges of M , so that M also has $O(g + b)$ edges.

Recall that every edge e of M belongs to a non-null-homotopic 3-cycle (Lemma 2); let C_e be such a cycle.

Construction of M_1 . Assume that there are three distinct edges e_1, e_2 , and e_3 of M such that C_{e_1} shares an edge with both C_{e_2} and C_{e_3} . Then $C_{e_1} \cup C_{e_2} \cup C_{e_3}$ has at most 5 vertices, implying that two of e_1, e_2, e_3 have an endpoint in common. This contradiction proves that every 3-cycle C_{e_1} shares an edge of T with at most one




Figure 5: In light lines, the matching M_2 ; in bold lines, the graph Π_2 , here forming a tree plus an edge.

other 3-cycle C_{e_2} . Let M_1 be obtained from M by removing e_1 or e_2 for every such pair of edges $\{e_1, e_2\}$. By the previous property, the cycles C_e , $e \in M_1$, are edge-disjoint. The set M_1 satisfies the hypotheses of the lemma, and $|M| \leq 2|M_1|$. Now, we forget M and focus on bounding the size of M_1 .

Construction of M_2 . We partition the edges e of M_1 according to the homotopy class of the corresponding 3-cycle C_e . Let M_2 be obtained by choosing one arbitrary representative edge per class; the cycles C_e , $e \in M_2$, are in distinct homotopy classes. We have $|M_1| \leq 9|M_2|$ by Lemma 6 and since the cycles C_e , $e \in M_1$, are edge-disjoint. Now, the cycles C_e , $e \in M_2$, are in distinct non-trivial homotopy classes and are edge-disjoint.

Construction of M_3 . For every $e \in M_2$, let π_e be the path of length two obtained from C_e by removing e . We orient the two edges of π_e towards the extremities of π_e (which are also the endpoints of $e \in M_2$). Since M_2 is a matching, every vertex of the triangulation T is the target of at most one oriented edge.

Let Π_2 be the union of the graphs π_e , $e \in M_2$. We claim that Π_2 is a *pseudoforest*: every connected component Π'_2 of Π_2 contains at most one cycle (see Figure 5). Indeed, every vertex of Π'_2 is the target of at most one oriented edge, so the number of edges of Π'_2 is at most the number of vertices of Π'_2 , and Π'_2 is connected; so a spanning tree of Π'_2 contains all but at most one edge of Π'_2 .

If Π'_2 is not a tree, let e' be an edge such that $\Pi'_2 - e'$ is a tree. The edge e' belongs to some π_e . We remove e from M_2 (and consequently π_e from Π'_2); the graph Π'_2 becomes one or two trees, and the other connected components of Π_2 are unaffected. We do this iteratively for every connected component Π'_2 of Π_2 . Let M_3 be obtained from M_2 after removing these edges.

Before the removal of any edge of M_2 , if a connected component Π'_2 of Π_2 is not a tree, it contains a cycle C of length at least three (Figure 5); since each vertex of C has at most one incoming edge, the edges of C belong to distinct C_e , for $e \in M_2$. Therefore, when removing edges of M_2 to form M_3 , we remove at most one third of the edges of M_2 . So $|M_2| \leq \frac{3}{2}|M_3|$.

End of the proof. Let Π_3 be the union of the graphs π_e , for $e \in M_3$; by construction, Π_3 is a forest. We now view T as a graph embedded on \mathcal{S} (slightly moving it towards the interior of \mathcal{S} , if \mathcal{S} has a boundary), and contract all the edges of Π_3 in this graph; this is legal since this set contains no cycle. Each edge e of M_3 is transformed into a loop ℓ_e homotopic to C_e . The loops ℓ_e form a graph Γ embedded on \mathcal{S} ; that graph has a single vertex per connected component. There exists a tree U embedded on \mathcal{S} meeting Γ exactly at its vertex set. We may contract U on the surface; now Γ is transformed into a set of simple, pairwise disjoint loops Γ' with the same vertex. Furthermore, the loops are non-null-homotopic and pairwise non-homotopic, so Corollary 5 implies that $|M_3| = |\Gamma'| \leq 3g + 2b - 3$ (unless \mathcal{S} is the projective plane, in which case the upper bound is one). By construction, we also have $|M| \leq 2|M_1| \leq 2 * 9|M_2| \leq 2 * 9 * \frac{3}{2}|M_3| = 27|M_3|$, which concludes the proof. \square

3.2 An Inclusionwise Maximal Matching Covers Many Vertices

Now, we prove that an inclusionwise maximal matching of T must cover a constant fraction of the edges of T .

Proposition 8. *Let T be an irreducible triangulation of \mathcal{S} . Let W be a set of vertices of T . Let M be an inclusionwise maximal matching of T among those that avoid W . Assume further that every boundary vertex of T is either in W or incident to an edge of M . Then the number of vertices of T is at most $7|M| + 4|W| + 3g + 3b - 6$.*

Proof. Let us denote by V , E , and F the vertices, edges, and faces of T , respectively. Let V_M be the vertices reached by M and X be the vertices neither in V_M nor in W . Let \bar{M} be the set of the edges of T that are not in M . Thus $\{W, V_M, X\}$ is a partition of V , and $\{M, \bar{M}\}$ is a partition of E .

Let $v \in X$. Recall that v is a non-boundary vertex by hypothesis. According to Lemma 3, v has degree at least four, so it is incident to at least four edges in \bar{M} . By the maximality of the matching M , the other vertex of each of these four edges is not in X . So, charging each vertex v of X with these four edges, we obtain that $4|X| \leq |\bar{M}|$.

The rest of the proof is standard machinery. Since T is a triangulation, by double-counting we obtain $|F| \leq \frac{2}{3}|E|$ (this is not an equality in general since \mathcal{S} may have boundary). Plugging this relation into Euler's formula $|V| - |E| + |F| = 2 - g - b$, we obtain:

$$(|W| + |V_M| + |X|) - \frac{1}{3}(|M| + |\bar{M}|) \geq 2 - g - b.$$

$|V_M| = 2|M|$ gives, after some rearranging:

$$|\bar{M}| - 3|X| \leq 5|M| + 3|W| + 3g + 3b - 6.$$

As shown above, $4|X| \leq |\bar{M}|$, implying $|X| \leq |\bar{M}| - 3|X|$ and so

$$|X| \leq 5|M| + 3|W| + 3g + 3b - 6.$$

This bound on $|X|$ allows to bound $|V| = 2|M| + |W| + |X|$ in terms of $|M|$, $|W|$, g , and b , implying the result. \square

3.3 End of Proof

The proof of Theorem 1 combines Propositions 7 and 8:

Proof of Theorem 1. Let W be a set of vertices, one on each boundary component of \mathcal{S} having an odd number of vertices. Build a matching M made of edges on the boundary of \mathcal{S} and covering the vertices on the boundary of \mathcal{S} that are not in W . Extend M to an inclusionwise maximal matching of T that avoids W ; we still denote it by M .

M contains no linking edge by construction so, by Proposition 7, M has less than $81g + 54b - 81$ edges (27 if \mathcal{S} is the projective plane). By Proposition 8, and since $|W| \leq b$, the number of vertices of T is at most $7|M| + 3g + 7b - 6$.

Combining these equations proves that T has at most $570g + 385b - 573$ vertices (186 if \mathcal{S} is the projective plane). \square

4 Improvement for Surfaces Without Boundary

The purpose of this section is to improve the previous bound when \mathcal{S} has no boundary ($b = 0$). The strategy is to improve the bound of Proposition 8 using a more careful analysis.

Theorem 9. *Let \mathcal{S} be a (possibly non-orientable) surface with Euler genus $g \geq 1$ and without boundary. Then every irreducible triangulation of \mathcal{S} has at most $f(g)$ vertices, where $f(1) = 55$, $f(2) = 194$, $f(3) = 333$, and $f(g) = 163g - 164$ if $g \geq 4$.*

The following lemma appears in an article by Fujisawa et al. [11, Sect. 2]; we reproduce the proof in more detail here for convenience.

Lemma 10. *Let \mathcal{S} be a surface of Euler genus $g \geq 1$ without boundary, and let $G = (V, E)$ be a 4-connected graph embedded on \mathcal{S} . Then for every $U \subseteq V$, the number of components of $G - U$ is at most $\max\{1, |U| + g - 2\}$.*

Proof. We can assume that $U \neq \emptyset$ and that $G - U$ has at least two connected components; otherwise, the result is clear. Let K be the graph obtained from G by the following steps:

1. Contract the edges of a spanning forest of $G - U$. Now the current graph has vertex set $U \cup W$, where W has one element for each component of $G - U$. The following steps will only add and remove edges of this graph.
2. Delete each edge with both endpoints in U . Similarly, delete each edge with both endpoints in $G - U$ (such edges are actually loops, by the first step). Now the current graph is bipartite.
3. On each face of the resulting graph that is not a disk, add edges to cut that face into a disk. This can be done without violating bipartiteness, because every face has a boundary component with at least one vertex in U and at least one vertex in W (since U and W are non-empty).
4. If there exists a face incident with exactly two edges, remove one of these two edges. (The two edges incident to the face are distinct, because \mathcal{S} is not the sphere and the edge is not a loop.) Repeat this step as much as possible.

We now have:

$$\begin{aligned} 4|W| &\leq |E(K)| \\ &\leq 2(|E(K)| - |F(K)|) \\ &= 2(|W| + |U| + g - 2). \end{aligned}$$

Indeed, the first inequality holds by 4-connectivity of G : since $|W| \geq 2$, every component of $G - U$ is adjacent to at least four different vertices of U ; therefore, in K , every vertex of W is adjacent to at least four different vertices of U . The second line follows from the fact that each face is incident to at least four edges (by bipartiteness of K and using Step 4). The third line holds by virtue of Euler's formula, since K is cellularly embedded on \mathcal{S} . \square

Proposition 11. *Let \mathcal{S} be a surface of Euler genus $g \geq 1$ without boundary, and let $G = (V, E)$ be a 4-connected graph embedded on \mathcal{S} . Let M be a maximum-size matching of G . Then the number of vertices of G is at most $2|M| + \max\{1, g - 2\}$.*

Proof. The Tutte-Berge formula [7, 36] [30, Sect. 24.1] asserts that the number of vertices of G not covered by a maximum-size matching of G is the maximum, over all $U \subseteq V$, of $o(G - U) - |U|$, where $o(G - U)$ denotes the number of components of the graph $G - U$ with an odd number of vertices. By Lemma 10, for every $U \subseteq V$, we have $o(G - U) - |U| \leq \max\{1, g - 2\}$. The result follows. \square

Proof of Theorem 9. If T is 4-connected, by Proposition 11, T has at most $2|M| + \max\{1, g - 2\}$ vertices where M is a maximum-size matching of T . Using the bound on the size of a maximal matching M (Proposition 7), we deduce that T has at most $h(g)$ vertices, where $h(1) = 55$, $h(2) = 163$, and $h(g) = 163g - 164$ if $g \geq 3$.

If T is not 4-connected, this means that a vertex set U of size at most three separates T . Actually, $|U| = 3$, and U forms a 3-cycle C in T . This cycle C must be separating, but also non-null-homotopic, for otherwise some edge of T would be contractible (as in the proof of Lemma 2). Let \mathcal{S}_1 and \mathcal{S}_2 be the surfaces obtained by cutting \mathcal{S} along C and attaching a triangle to each copy of C . The Euler genera of \mathcal{S}_1 and \mathcal{S}_2 add up to g . Furthermore, C is two-sided (since it is separating), so the number of 3-cycles homotopic to C in \mathcal{S} is at most 27 [5, Lemma 9]. Any edge that is contractible in \mathcal{S}_1 or \mathcal{S}_2 belongs to such a cycle. So the total number of edges in \mathcal{S}_1 and \mathcal{S}_2 that are contractible is at most $3 \times 27 + 3 = 84$ (the “+3” term comes from the fact that the three edges of C may be contractible in both \mathcal{S}_1 and \mathcal{S}_2 .) A similar reasoning is used by Barnette and Edelson [5, Proof of Theorem 2].

It follows that the number of vertices of an irreducible triangulation of a surface without boundary with Euler genus g is bounded from above by $f(g)$, where f satisfies the induction formula:

$$f(g) = \max \left\{ h(g), \max_{\substack{g_1+g_2=g \\ g_1, g_2 \geq 1}} \{f(g_1) + f(g_2) + 84\} \right\}.$$

Thus, we have $f(1) = 55$, $f(2) = 194$, and for $g \geq 3$:

$$f(g) = \max \left\{ 163g - 164, \max_{\substack{g_1+g_2=g \\ g_1, g_2 \geq 1}} \{f(g_1) + f(g_2) + 84\} \right\}.$$

It is easily checked by induction that $f(3) = 333$ and $f(g) = 163g - 164$ for $g \geq 4$. □

Acknowledgment. We would like to thank Gwenaël Joret, Thom Sulanke, and David Wood for helpful comments which improved the readability of this article.

References

- [1] Dan Archdeacon. The nonorientable genus is additive. *Journal of Graph Theory*, 10(3):363–383, 1986.
- [2] Mark Anthony Armstrong. *Basic topology*. Undergraduate Texts in Mathematics. Springer-Verlag, 1983.
- [3] Dominique Attali, André Lieutier, and David Salinas. Efficient data structure for representing and simplifying simplicial complexes in high dimensions. In *Proceedings of the 27th Annual Symposium on Computational Geometry (SOCG)*, pages 501–509. ACM, 2011.
- [4] David Barnette. Generating the triangulations of the projective plane. *Journal of Combinatorial Theory, Series B*, 33:222–230, 1982.
- [5] David W. Barnette and Allan Edelson. All orientable 2–manifolds have finitely many minimal triangulations. *Israel Journal of Mathematics*, 62:90–98, 1988.
- [6] David W. Barnette and Allan Edelson. All 2–manifolds have finitely many minimal triangulations. *Israel Journal of Mathematics*, 67:123–128, 1989.
- [7] Claude Berge. Sur le couplage maximum d’un graphe. *Comptes rendus hebdomadaires des séances de l’Académie des Sciences [Paris]*, 247:258–259, 1958.
- [8] Erin W. Chambers, Éric Colin de Verdière, Jeff Erickson, Francis Lazarus, and Kim Whittlesey. Splitting (complicated) surfaces is hard. *Computational Geometry: Theory and Applications*, 41(1–2):94–110, 2008.
- [9] Siu-Wing Cheng, Tamal K. Dey, and Sheung-Hung Poon. Hierarchy of surface models and irreducible triangulations. *Computational Geometry: Theory and Applications*, 27(2):135–150, 2004.
- [10] Tamal K. Dey, Herbert Edelsbrunner, Sumanta Guha, and Dmitry V. Nekhayev. Topology preserving edge contraction. *Publications de l’Institut Mathématique (Beograd) (N.S.)*, 66(80):23–45, 1999. Geometric combinatorics (Kotor, 1998).
- [11] Jun Fujisawa, Atsuhiko Nakamoto, and Kenta Ozeki. Hamiltonian cycles in bipartite toroidal graphs with a partite set of degree four vertices. *Journal of Combinatorial Theory, Series B*, 103(1):46–60, 2013.
- [12] Zhicheng Gao, L. Bruce Richmond, and Carsten Thomassen. Irreducible triangulations and triangular embeddings on a surface. Technical Report CORR 91-07, University of Waterloo, Canada, 1991. Available at <http://mathstat.carleton.ca/~zgao/PAPER/GRT91.pdf>.

- [13] Zhicheng Gao, R. Bruce Richter, and Paul D. Seymour. Irreducible triangulations of surfaces. *Journal of Combinatorial Theory, Series B*, 68:206–217, 1996.
- [14] Michael Henle. *A combinatorial introduction to topology*. Dover Publications, 1994.
- [15] Gwenaël Joret and David R. Wood. Irreducible triangulations are small. *Journal of Combinatorial Theory, Series B*, 100:446–455, 2010.
- [16] Mark Jungerman and Gerhard Ringel. Minimal triangulations on orientable surfaces. *Acta Mathematica*, 145(1):121–154, 1980.
- [17] Martin Juvan, Aleksander Malnič, and Bojan Mohar. Systems of curves on surfaces. *Journal of Combinatorial Theory, Series B*, 68:7–22, 1996.
- [18] Ken-ichi Kawarabayashi, Raiji Mukae, and Atsuhiro Nakamoto. K_6 -minors in triangulations on the Klein bottle. *SIAM Journal on Discrete Mathematics*, 23:96–108, 2008.
- [19] Serge Lawrencenko. The irreducible triangulations of the torus. *Ukrainskiĭ Geometricheskiiĭ Sbornik*, 30:52–62, 1987.
- [20] Aleksander Malnič and Roman Nedela. k -minimal triangulations of surfaces. *Acta Mathematica Universitatis Comenianae*, 64(1):57–76, 1995.
- [21] Gary L. Miller. An additivity theorem for the genus of a graph. *Journal of Combinatorial Theory, Series B*, 43(1):25–47, 1987.
- [22] Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 2001.
- [23] Raiji Mukae and Atsuhiro Nakamoto. K_6 -minors in triangulations on the nonorientable surface of genus 4. In preparation.
- [24] Raiji Mukae and Atsuhiro Nakamoto. K_6 -minors in triangulations and complete quadrangulations. *Journal of Graph Theory*, 60:302–312, 2009.
- [25] Raiji Mukae, Atsuhiro Nakamoto, Yoshiaki Oda, and Yusuke Suzuki. K_6 -minors in triangulations on the nonorientable surface of genus 3. *Graphs and Combinatorics*, 26:559–570, 2010.
- [26] Atsuhiro Nakamoto and Katsuhiko Ota. Note on irreducible triangulations of surfaces. *Journal of Graph Theory*, 20(2):227–233, 1995.
- [27] Seiya Negami. Diagonal flips of triangulations on surfaces, a survey. *Yokohama Mathematical Journal*, 47:1–40, 1999.
- [28] Seiya Negami. Diagonal flips in triangulations of surfaces. *Discrete Mathematics*, 135(1–3):225–232, 1994.
- [29] Haijo Schipper. Generating triangulations of 2-manifolds. In *Computational Geometry—Methods, Algorithms and Applications: Proc. Internat. Workshop Comput. Geom. CG '91*, volume 553 of *Lecture Notes in Computer Science*, pages 237–248. Springer-Verlag, 1991.
- [30] Alexander Schrijver. *Combinatorial optimization. Polyhedra and efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003.
- [31] Ernst Steinitz and Hans Rademacher. *Vorlesungen über die Theorie der Polyeder*. Springer-Verlag, 1934.
- [32] John Stillwell. *Classical topology and combinatorial group theory*. Springer-Verlag, New York, second edition, 1993.
- [33] Thom Sulanke. Generating irreducible triangulations of surfaces. arXiv:math/0606687, 2006.

- [34] Thom Sulanke. Irreducible triangulations of low genus surfaces. arXiv:math/0606690, 2006.
- [35] Thom Sulanke. Note on the irreducible triangulations of the Klein bottle. *Journal of Combinatorial Theory, Series B*, 96:964–972, 2006.
- [36] Douglas B. West. A short proof of the Berge–Tutte formula and the Gallai–Edmonds structure theorem. *European Journal of Combinatorics*, 32(5):674–676, 2011.