The Quadrangulation Conjecture for Orientable Surfaces

Daniel R. L. Brown*
Certicom Research
5520 Explorer Drive, 4th Floor, Mississauga, Ontario, Canada, L4W 5L1
E-mail: dbrown@certicom.com

David M. Jackson†
Department of Combinatorics and Optimization
University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1
E-mail: dmjackson@math.uwaterloo.ca

December 6, 2001

Abstract

By means of character theory and symmetric functions, Jackson and Visentin [4] proved the existence of certain bijections between the set of quadrangulations in orientable surfaces and decorated maps (with marked edges and coloured vertices) in orientable surfaces. The bijections preserve a weight function consisting of a pair (g, n) of integer parameters. For quadrangulations, g is the genus and n is the number of faces. For decorated maps, g is the genus plus half the number of white vertices and n is the number of edges.

The Quadrangulation Conjecture concerns the problem of finding a natural bijection of this type. Tutte's medial construction is a solution in the special case $g = 0$ of planar maps. We give a construction of a bijection $\widetilde{\Xi}$ which both extends Tutte's medial construction to non-planar maps and preserves the parameter n of the Quadrangulation Conjecture. (The parameter g is not generally preserved, except when $g = 0$.) Non-orientable surfaces play an important part in the construction of $\widetilde{\Xi}$. As part of the construction, we introduce a bijection between orientable rooted quadrangulations and locally orientable, bipartite, rooted quadrangulations.

1 Introduction

A *map* is an embedding of a graph, with vertices sent to points and edges sent to continuous paths (joining the images of the vertices to which the edge is incident), into a closed surface without boundary (compact 2-manifold), such that the remainder of the surface has components each of which is homeomorphic to an open disk. For fixed values of certain sets of parameters of maps, the number of maps is, up to isomorphism, finite. The specific topic of this paper is a problem from the enumeration of maps in surfaces, the Quadrangulation Conjecture [4], involving the construction of a natural bijection between rooted quadrangulations and decorated rooted maps in arbitrary orientable surfaces that preserves parameters g and n (where g is the genus of a quadrangulation and also the genus plus half the number of white-coloured vertices of a decorated map, n is the number of faces of a quadrangulation and also the number of edges of a decorated map). The existence of a bijection between these sets is assured by an enumerative result in [4] relating the generating series of the two sets. However, the proof contained algebraic steps which defy direct combinatorial interpretation and therefore offered very little help in the determination of a natural bijection.

The approach taken in this paper is to use depth first search to label rooted maps canonically and naturally, thus facilitating further combinatorics on rooted maps. With such methods, a new advance is

*Supported by the National Sciences and Engineering Research Council of Canada.
†Supported by the National Sciences and Engineering Research Council of Canada.

made which partially solves the Quadrangulation Conjecture. The approach uses both orientable and non-orientable surfaces. As part of the bijection for the Quadrangulation Conjecture, we give an intermediate bijection between orientable quadrangulations and bipartable locally orientable quadrangulations.

The conjecture and its enumerative background. We begin with a brief account of the Quadrangulation Conjecture and its enumerative background. Let $M(u, x, y, z)$ and $Q(u, x, y, z)$ be the generating series for rooted orientable maps and rooted quadrangulations (face-4-regular maps) respectively, where, in the generating series the indeterminates u, x, y, z mark genus, faces, vertices and edges, respectively. (A map is rooted by selecting an end of an edge, and a side of that end. See §2 for greater detail.) The indeterminate u is redundant by the Euler-Poincaré formula, but is included for convenience. A factorisation result [4] leads to the following remarkable relationship:

$$Q(u^2, x, y, z) = \mathsf{bis}_u M(4u^2, y + u, y, xz^2), \tag{1}$$

where $\mathsf{bis}_u f$ denotes the even bisection $\frac{1}{2}\{f(u) + f(-u)\}$ of the formal power series f.

The result (1) implies existence of bijections between two sets of rooted maps, namely $\mathcal{Q}$, the set of rooted orientable quadrangulations, and $\mathcal{A}$, the set of decorated rooted orientable maps. Decorated maps are defined in §4 and examples are given in Table 1(b). Decorations are introduced to account for the right-hand side of (1). There are pairs (g, n) of corresponding weight-functions defined on $\mathcal{Q}$ and $\mathcal{A}$, and the generating series equality (1) further implies that there exists a weight-preserving bijection $\xi : \mathcal{Q} \to \mathcal{A}$.

Conjecture 1 (Quadrangulation Conjecture [4]) *There exists a (g, n)-preserving bijection $\Xi : \mathcal{Q} \to \mathcal{A}$ admitting an element-wise action having a combinatorial description.*

The construction in Conjecture 1 for the bijection Ξ, if found, would provide a new and purely combinatorial proof of (1). In §5, a bijection $\widetilde{\Xi} : \mathcal{Q} \to \mathcal{A}$ which preserves one of the two required weight functions is found. Thus $\widetilde{\Xi}$ is a partial solution to Conjecture 1.

There are several implications of the combinatorial proof to the conjecture. For example, the enumerative result (1) proved [3] physicists' suspected connection between the φ^4-model and Penner model of 2-dimensional quantum gravity (see also [7]). The related hypothetical bijection Ξ of the Quadrangulation Conjecture could provide more details of a physical interpretation of the connection between these two models, at the level of Feynman diagrams.

The following theorem is the main result. It first appeared in [1] and partially establishes the Quadrangulation Conjecture.

Theorem 1 *There exists a n-preserving bijection $\widetilde{\Xi} : \mathcal{Q} \to \mathcal{A}$ extending Tutte's medial bijection and admitting an element-wise action having a combinatorial description.*

The element-wise action of $\widetilde{\Xi}$ is specified in §5 proving Theorem 1. Since $\widetilde{\Xi}$ extends Tutte's medial bijection, which is a solution of the Quadrangulation Conjecture for the special case of $g = 0$, the bijection $\widetilde{\Xi}$ preserves g if $g = 0$.

2 Maps and Matchings

A *locally orientable surface* [10] is a compact 2-manifold. Up to homeomorphism, two parameters, the Euler characteristic $\chi \in \mathbb{Z}$ and orientability, suffice to determine locally orientable surfaces. Orientable surfaces are the sphere with g handles, with Euler characteristics $\chi = 2 - 2g$, where g is the *genus* of the orientable surface. Non-orientable surfaces are the real projective plane or Klein bottle with g handles, with Euler characteristic $\chi = 1 - 2g$ or $\chi = -2g$, respectively. A topological map is an ordered triple $\mathfrak{M} = (G, S, i)$ where i is a *2-cell embedding* of a graph G in a surface S. The 2-cells are the *faces* of the map $\mathfrak{M}$. The Euler characteristic of S may be computed as $\chi = f - e + v$, where f, e, v are the number of faces, edges and vertices of a topological map $\mathfrak{M} = (G, S, i)$. Topological maps $\mathfrak{M} = (G, S, i)$ and $\mathfrak{M}' = (G', S', i')$ are isomorphic if there is a homeomorphism $h : S \to S'$ which preserves the images of the graphs, *i.e.* $h(i(G)) = i'(G')$. It follows that if $\mathfrak{M}$ and $\mathfrak{M}'$ are isomorphic, $\mathfrak{M} \cong \mathfrak{M}'$, then the underlying graphs are isomorphic, $G \cong G'$ and the underlying surfaces are homeomorphic, $S \cong S'$. (The converse,

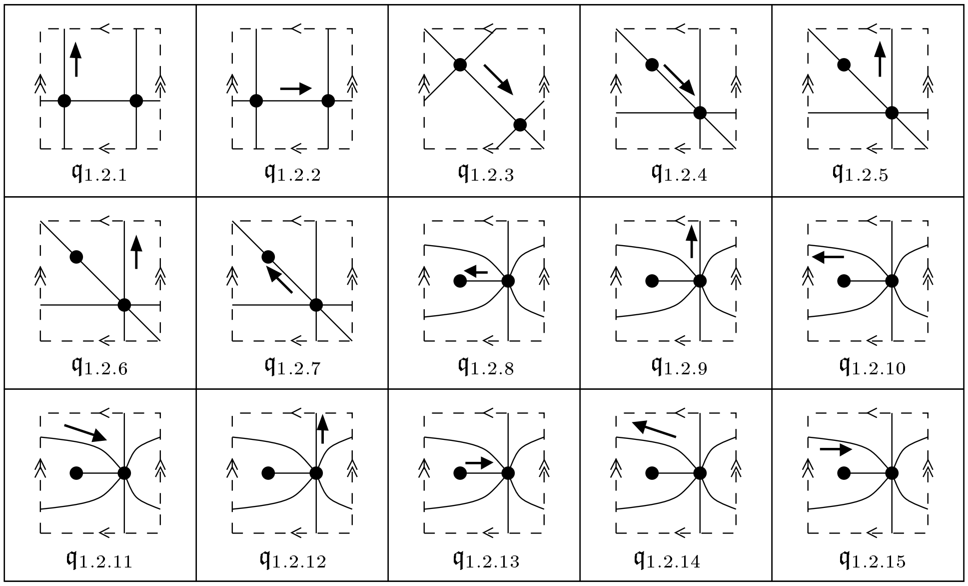

(a) $\mathcal{Q}_{1,2}$ —Two-faced quadrangulations of genus one

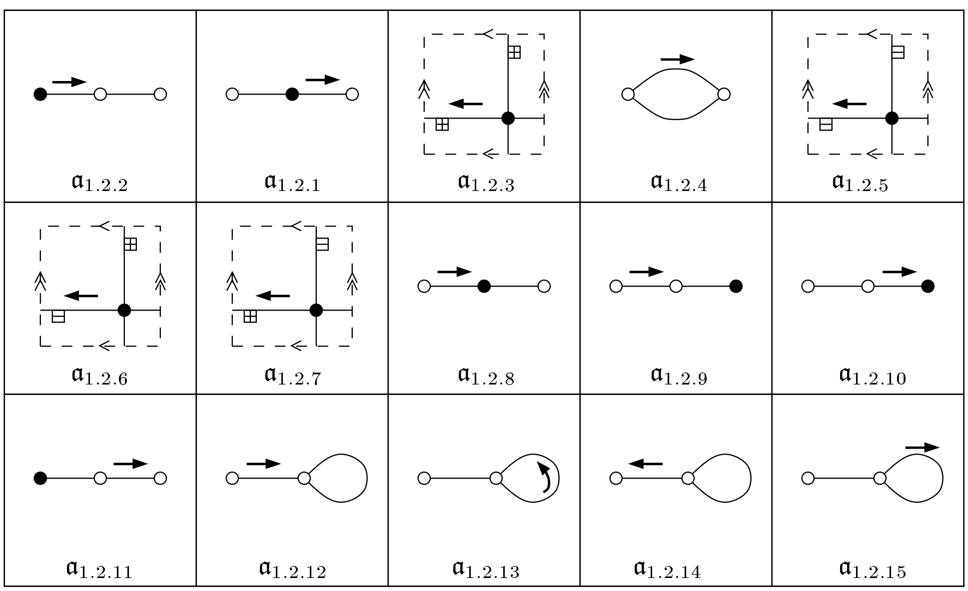

(b) $\mathcal{A}_{1,2}$ — Decorated, two-edged maps (with $w + h = 1$)

Table 1: Listing of the element-wise actions $\widetilde{\Xi}\mathfrak{q}_{1.2.i} = \mathfrak{a}_{1.2.i}$ for $i = 1, \ldots, 15$ (see §4 for notation)

however, need not be true. For example, the graph consisting of a two edge path plus a loop incident to the middle vertex of the path, has two non-isomorphic embeddings in the sphere. On the other hand, 3-connected graphs [10] have at most one embedding in the sphere.)

For combinatorial purposes it is convenient to have a purely combinatorial definition. The set of permutations π of a finite set X is denoted by $\mathrm{Sym}(X) = \{\pi : X \to X \mid \pi \text{ is bijective}\}$. Let $\mathfrak{S}_n = \mathrm{Sym}(\{1, 2, \ldots, n\})$. The lengths of the cycles of a permutation π form a partition $\tau(\pi)$, its *cycle type*. A *matching* is a permutation μ whose cycle type is of the form $\tau(\mu) = [2^n]$. The set of matchings on a set X is $\mathrm{Match}(X) \subset \mathrm{Sym}(X)$. The following definition is from [2] and was adapted from a definition of Tutte [10].

Definition 1 *A quadruple* $\mathfrak{m} = (X, \mu_1, \mu_2, \mu_3)$ *is a* combinatorial map *if:*

1. *X is a finite set,*
2. *μ_k is a matching on X, for $k = 1, 2, 3$,*
3. *The permutation product $\mu_1\mu_3$ is also a matching on X,*
4. *The group $\langle \mu_1, \mu_2, \mu_3 \rangle$ acts transitively on X.*

Combinatorial maps $\mathfrak{m} = (X, \mu_1, \mu_2, \mu_3)$ and $\mathfrak{m}' = (X', \mu_1', \mu_2', \mu_3')$ are isomorphic if and only if there is a bijection $\beta : X \to X'$ such that $\beta \circ \mu_k = \mu_k' \circ \beta$ for $k = 1, 2, 3$. The following transformation takes a topological map $\mathfrak{M} = (G, S, i)$ and produces a combinatorial map and gives, up to isomorphism, a one-to-one correspondence between topological and combinatorial maps. Therefore the problems of enumerating combinatorial maps and enumerating topological maps are equivalent. For each edge $e \in G$, the curve $i(e)$ has two sides and two ends. Choose four points on S, each very close to one side and one end of $i(e)$. Having done this for each edge e, forming a set of such points X, let these points be the *(side-end) positions* of $\mathfrak{M}$. Thus, if G has n edges, then of course $|X| = 4n$. The side-end positions are related to each other by the local topology of $\mathfrak{M}$. More precisely, the positions in X can be naturally paired in three different ways: pairs in X associated with the same end of one edge but on different sides giving the *same-end matching* μ_1; pairs in X belonging to the same corner of a face giving the *same-corner matching* μ_2; and pairs in X associated with the same side of one edge, but at different ends giving the *same-side matching* μ_3. For example, in Figure 7 (left), positions c and C are at the 'same-end', c and B are at the 'same-corner', and c and E are at the 'same-side'.

The *matchings graph* of $\mathfrak{m}$ is the graph with vertex set X and edge set given by the matchings μ_i. The edges of $\mathfrak{m}$ correspond to 4-cycles of the matchings graph formed by the matchings μ_1 and μ_3. The dual of the map $\mathfrak{m}$ is $\mathfrak{m}^\partial = (X, \mu_3, \mu_2, \mu_1)$. Topologically, in the dual each vertex is replaced by a face, each face by a vertex. For each edge e of $\mathfrak{m}$ there is an associated edge e^∂ of $\mathfrak{m}^\partial$.

Definition 1 includes non-orientable maps. Although the Quadrangulation Conjecture only refers to orientable maps, non-orientable maps will play an important part in the construction of $\widetilde{\Xi}$. A *rooted map* is a quintuple $\mathfrak{m} = (X, \mu_1, \mu_2, \mu_3, r)$ where $r \in X$ and (X, μ_1, μ_2, μ_3) is a combinatorial map. The Quadrangulation Conjecture is a statement concerning rooted maps. The position r is the *root* and plays an important part in the construction of $\widetilde{\Xi}$.

3 Canonical Representation of Rooted Maps

Tutte noted in [10] that maps are more easily enumerated when rooted. A reason for this is that rooting a map removes all automorphisms but the identity. There are two immediate observations about rooting. First, combinatorial decompositions may be found for rooted maps, and second, there is a very simple relation between numbers of labelled maps and numbers of rooted maps. The first observation was used by Tutte and others for very direct enumeration of rooted maps, typically in the form of a decomposition that operates locally at the root of the map (for example, by root edge deletion). The second observation is used in the algebraic method of map enumeration, and has lead to results such as [4, 2]. A direct combinatorial interpretation for rooted maps of some of the enumerative results has not been found, the complexity of algebra involved in the proofs being an obstruction to such an interpretation (combinatorial interpretations have recently been found [8] for some of Tutte's results [9]). A second obstruction is that the algebraic proofs are based on the enumeration of labelled maps rather than rooted maps.

These obstructions are avoided by seeking more complex direct decompositions of rooted maps. In particular, because rooted maps have no automorphism, and thus no parts of a rooted map are equivalent, it follows that decompositions of rooted maps can be defined that depend globally on the rooted map. For example, Walsh and Lehman [5, 6, 11] have used such decompositions. This section extends these decompositions and derives further ones, as preparation for the definition of a bijection $\tilde{\Xi}$ that is a partial solution to the Quadrangulation Conjecture.

3.1 The Canonical Position Labelling Algorithm

The following depth first search algorithm defines a canonical labelling $\ell_{\mathfrak{m}} : X \to \mathcal{N}_{|X|}$ of the position set of an arbitrary locally orientable, rooted map $\mathfrak{m} = (X, \mu_1, \mu_2, \mu_3, a)$, where $\mathcal{N}_i = \{1, 2, \ldots, i\}$ for i a positive integer.

Algorithm 1 (Canonical Position Labelling) *Let* $\mathfrak{m} = (X, \mu_1, \mu_2, \mu_3, a)$ *be a rooted map.*

1. *Let* $j := 1$ *and* $x := a$.
2. *Let* $x_j = x$.
3. *Let* $k := 1$.
4. *If* $\mu_k(x) \notin \{x_1, \ldots, x_j\}$ *then let Parent*$(\mu_k(x)) = x$, $x := \mu_k(x)$, $j := j + 1$ *and goto Step 2.*
5. *Let* $k := k + 1$. *If* $k \leqslant 3$ *then goto Step 4.*
6. *If* $x \neq a$ *then let* $x :=$ *Parent*(x) *and goto Step 3.*
7. *Halt.*

Then $\ell_{\mathfrak{m}} : X \to \mathcal{N}_{|X|} : x_j \mapsto j$, *is called the* canonical position labelling *of* $\mathfrak{m}$.

Because this is a depth first search algorithm in the sense of the matchings graph of $\mathfrak{m}$, it will halt having labelled each of the positions in X to produce a unique labelling for any rooted map. Thus each rooted map is isomorphic to exactly one canonical map as defined below.

Definition 2 *A rooted map* $\mathfrak{m} = (\mathcal{N}_{4n}, \mu_1, \mu_2, \mu_3, 1)$ *is* canonical *if* $\ell_{\mathfrak{m}}(i) = i$ *for all* $1 \leqslant i \leqslant 4n$.

In the examples it will be helpful to use letters $a, b, c, \ldots$ or $a, A, b, B, \ldots$ for the canonical position labels $1, 2, 3, \ldots$ (in order to the avoid the use of double digit labels). The letters e, f, g, v, however, retain their original definitions as map invariants. Which of the two meanings is appropriate will be clear from the context.

Example of the execution of Algorithm 1. Consider the planar rooted map given on the left in Figure 1. The tail of the arrow in Figure 1 indicates the root position of this map, which is a. The root position is the starting point for the algorithm. In the middle diagram, labels $a, b, c, d, \ldots$ have been added, according to Algorithm 1, such that $\ell_{\mathfrak{m}}(a) = 1, \ell_{\mathfrak{m}}(b) = 2, \ell_{\mathfrak{m}}(c) = 3, \ldots$. The trace of the algorithm is indicated by the curve superimposed on the map on the right in Figure 1. Such curves can in general be drawn to trace the algorithm, without intersecting themselves. Note that Algorithm 1 backtracks in the example from r to q to p before assigning label s. In Algorithm 1, backtracking occurs in Step 6 if all positions adjacent to the current position have previously been labelled, as is the case for position r in the example whose neighbouring positions were previously labelled q, m, a.

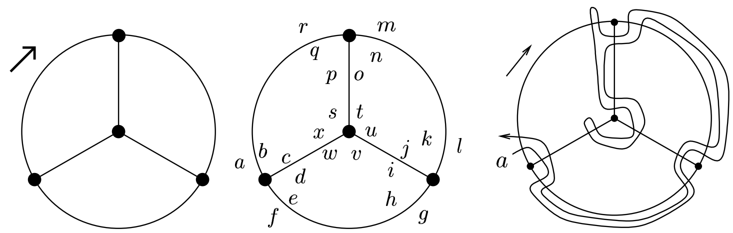

Figure 1: A planar rooted embedding of the complete graph K_4, its canonical position labelling and the trace of the Algorithm 1.

3.2 Canonical Edge Ordering

The edges have a natural *canonical order*: sorted by ascending order of their highest canonical position label. In the example of Figure 1, the sets of edge labels are $\{e, f, g, h\}, \ldots, \{c, d, w, x\}$, so the assignments of the canonical edge labels $e_1, \ldots, e_6$ to the edges are

$$\begin{array}{ccc} e_1 \leftarrow \{e, f, g, h\}, & e_2 \leftarrow \{k, l, m, n\}, & e_3 \leftarrow \{a, b, q, r\}, \\ e_4 \leftarrow \{o, p, s, t\}, & e_5 \leftarrow \{i, j, u, v\}, & e_6 \leftarrow \{c, d, w, x\}. \end{array}$$

The canonical edge ordering is $e_1, \ldots, e_6$.

3.3 Canonical Spanning Trees, Struts and Cuts

Usually, for the occurrences of Step 4 of Algorithm 1 in which the re-assignment $x := \mu_k(x)$ is made, the value of k is 1 or 2, because these values are considered before the value $k = 3$. In fact, $k = 3$ for $v - 1$ of these re-assignments, because one such re-assignment $x := \mu_3(x)$ is needed each time the algorithm visits a new vertex (excluding the root vertex). When this happens, we say that the edge of $\mathfrak{m}$ containing the positions x and $\mu_3(x)$ has been *traversed* by Algorithm 1. The set of the traversed edges of the rooted map $\mathfrak{m}$ form a spanning subtree of the underlying graph of $\mathfrak{m}$. (Not the matchings graph.) This tree is the *canonical spanning subtree* of $\mathfrak{m}$. Its edges are called *canonical tree edges.* For example, in Figure 1, the tree edges are the three edges having labels $efgh$, $klmn$, and $opst$.

Cuts and Struts. Depth first search has been used to classify edges canonically into tree edges and non-tree edges. Later, in the construction of the bijection $\widetilde{\Xi}$, we shall need the following more refined classification, one which separates non-tree edges into *cuts* and *struts.*

Algorithm 2 (Cuts and Struts) *Let $\mathfrak{m}$ be a rooted map with n edges. Determine the canonical tree edges of $\mathfrak{m}$. Then*

1. *Let $i := n$ and $\mathfrak{b} := \mathfrak{m}$, $C_{\mathfrak{m}} := \varnothing$. Let $e_1, e_2, \ldots$ be the canonical edge ordering.*
2. *If e_i is not a canonical tree edge and bounds two distinct faces in $\mathfrak{b}$, then let $C_{\mathfrak{m}} := C_{\mathfrak{m}} \cup \{e_i\}$ and let $\mathfrak{b} := \mathfrak{b} - e_i$.*
3. *Let $i := i - 1$. If $i > 0$, goto Step 2.*
4. *Halt.*

Then the edges in $C_{\mathfrak{m}}$ are called the cuts *of $\mathfrak{m}$, the terminal value of the submap $\mathfrak{b}$ is the* backbone *of $\mathfrak{m}$, and the edges of $\mathfrak{b}$ which are not canonical tree edges are called the* struts *of $\mathfrak{m}$.*

There are $v - 1$ tree edges. Because the number of faces is reduced by one each time a cut is removed from $\mathfrak{b}$, there are $f - 1$ cuts. Therefore there are $2 - \chi$ struts. For orientable maps, $\chi = 2 - 2g$, so the number of struts is $2g$.

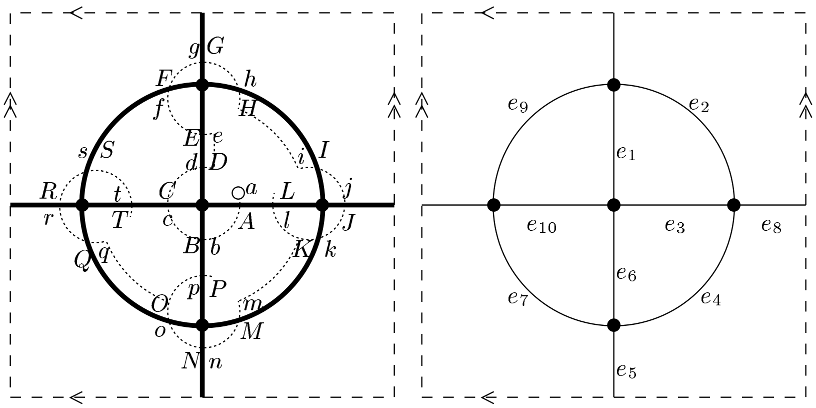

Figure 2: The canonical position labelling and edge ordering for a K_5 in the torus

Example of the determination of cuts and struts. Consider an example of the rooted embedding $\mathfrak{m}$ of K_5 in the torus, given in Figure 2. First, perform Algorithm 1 on the positions of the map $\mathfrak{m}$, as in the left of Figure 2, for mnemonic ease the same-end position labels we use are a, A; b, B, *etc.* From §3.2, edges have a canonical order based on the order of the highest letter (position) on each edge. For $\mathfrak{m}$, this edge order is indicated in the right of Figure 2 by the symbols $e_1, e_2, \ldots, e_{10}$ in increasing order of the subscripts. Edge order plays a prominent role in distinguishing struts from cuts.

As before, certain edges are canonical tree edges and these form a canonical spanning subtree of the underlying graph K_5. In $\mathfrak{m}$, the tree edges are e_1, e_2, e_4, e_7, and are shown with thickened line segments in Figure 3. The non-tree edges of $\mathfrak{m}$ in descending order are $e_{10}, e_9, e_8, e_6, e_5, e_3$. To test which of these are struts and which are cuts, delete them in succession by descending order. Delete an edge only if it separates two distinct faces. The edges that are deleted in this process are the *cuts*. The edges that are not deleted are the *struts*. The result of this deletion process is, in general, a submap of $\mathfrak{m}$, that consists of only the tree edges and struts, has only one face, and has the same surface structure as $\mathfrak{m}$. This submap is the backbone of the rooted map $\mathfrak{m}$ and is given in Figure 4.

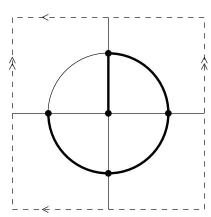

Figure 3: The canonical tree edges

Since e_{10} separates two distinct faces in $\mathfrak{m}$, e_{10} is a cut. Therefore, delete e_{10} from $\mathfrak{m}$. Next, e_9 separates two distinct faces in $\mathfrak{m} - e_{10}$. Therefore, delete e_9 too. Thus e_9 is a cut. In $\mathfrak{m} - \{e_{10}, e_9\}$, the next edge to test, e_8, does not separate two faces. Therefore e_8 is a strut. It is not deleted by the process. To emphasise that e_8 is a strut, it is now drawn with a double line, as in Figure 4. The next non-tree edge is e_6 which is another cut. Finally, e_5 is a strut and e_3 is a cut, for reasons similar to the previous.

Associating Cuts with Non-root Faces. In the example of K_5 embedded in the torus (see Figure 2), the backbone consists of the edges $\{e_1, e_2, e_4, e_5, e_7, e_8\}$, and is given in Figure 4. The missing edges are cuts. Form submaps $\mathfrak{m}_i$ of $\mathfrak{m}$ by taking the backbone to be $\mathfrak{m}_1$ and successively adding the cuts, in *ascending* order. The resulting maps are $\mathfrak{m}_1, \mathfrak{m}_2, \ldots, \mathfrak{m}_f = \mathfrak{m}$, where $\mathfrak{m}$ has f faces. Then $\mathfrak{m}_i$ in this sequence has i faces. As each cut is added, it separates an existing face into two faces, one of which will be considered to

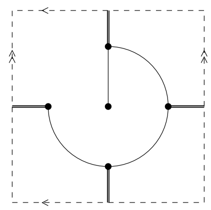

Figure 4: The backbone of the rooted embedding of K_5

be the *new* face, in a manner to be described shortly. This new face will become associated with the cut. Of course, subsequently, the face of $\mathfrak{m}_{i+1}$ that is associated with the i^{th} cut, may itself be divided into faces by later cuts. Nevertheless, the association remains between the original cut, and whatever remains after new faces are cut off the face associated with the original cut.

The root face is treated differently. The face R_i of $\mathfrak{m}_i$ containing the root position of $\mathfrak{m}$ is the one face of $\mathfrak{m}_i$ that is not associated with any cuts. If the next cut c separates R_i into two faces R_{i+1} and S_{i+1} where R_{i+1} is the face of $\mathfrak{m}_{i+1}$ that contains the root position of $\mathfrak{m}$, then the face S_{i+1} is associated with c. Thus, R_{i+1} is again the face of $\mathfrak{m} + i + 1$ that both contains the root position of $\mathfrak{m}$ and is the only face not associated with any cuts. Eventually all the non-root faces of $\mathfrak{m}$ will be associated with cuts.

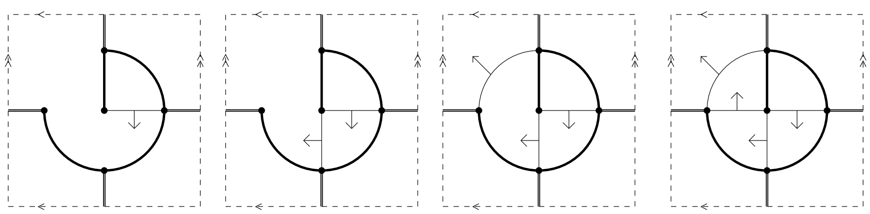

Figure 5: The submaps $\mathfrak{m}_2, \mathfrak{m}_3, \mathfrak{m}_4$ and $\mathfrak{m}_5$, with associations of e_3, e_6, e_9 and e_{10} to non-root faces

With the example, $\mathfrak{m}_2 = \mathfrak{m}_1 + e_3$, the backbone with e_3 (the lowest cut) added, as in Figure 5, on the left. Here, e_3 cuts the root face of $\mathfrak{m}_1$ into two faces. Since the root position of $\mathfrak{m}$ actually lies on e_3 itself, it is straightforward to see which of the two faces in e_3 is the non-root face. To indicate that the cut e_3 is associated with the non-root face of $\mathfrak{m}_2$, an arrow is drawn in the right of Figure 5, extending from one side of e_3 into the face with which it is associated. The next cut, e_6, cuts the face of $\mathfrak{m}_2$ associated with e_3. In $\mathfrak{m}_3$ one of these two resulting faces is incident with e_3 and the other not. The other face is the *new* face. It is not a root face and is not yet associated with any cuts, and we therefore associate it with the cut e_6. This is indicated in Figure 5, with an arrow on the side of e_6. The next two cuts, e_9 and e_{10}, also cut non-root faces, so are treated similarly to e_6.

3.4 Canonical Edge Diagram Structure

The canonical position labelling algorithm (Algorithm 1) sorts the positions of a rooted map into linear order. The *edge diagram* is a linearisation of a rooted map, in the sense that the positions are ordered linearly on a base line.

Algorithm 3 (ε — Canonical Edge Diagram) *Let $\mathfrak{m} = (X, \mu_1, \mu_2, \mu_3, a)$ be a rooted map with n edges and v vertices. Let $\ell_{\mathfrak{m}}$ be its canonical position labelling computed by Algorithm 1. The* (canonical) edge diagram $\mathfrak{e} = \varepsilon(\mathfrak{m})$ *of $\mathfrak{m}$ is formed of the following:*

1. *A horizontal line segment, whose left endpoint is P_0.*
2. *n semicircular arcs, where the interior of each arc are positioned above the line segment, and also the endpoints of the arcs intersect the interior line segment at $2n$ distinct points, $P_1, \ldots, P_{2n}$ from*

left to right such that P_i and P_j are the two endpoints of a common arc if and only if the positions of $\mathfrak{m}$ with canonical labels $2i$ and $2j$ belong to a common edge.

3. *v mutually disjoint small discs with centres positioned at P_0 and each point P_i where $i \neq 1$ and $2i-1$ is smallest the canonical label of the positions belonging to a vertex of $\mathfrak{m}$. The radii are small enough that each disc contains just one of the points P_i.*

4. *Some number of $\times$ symbols centred at points P_j such that P_j is on an arc with P_i where $i < j$ and $\mu_3(x) = \mu_3(y)$ where x and y are the unique positions of $\mathfrak{m}$ such that $\ell_{\mathfrak{m}}(x) = 2i$ and $\ell_{\mathfrak{m}}(y) = 2j$ respectively. (The symbol meets the diagram only at P_j.)*

In view of the construction of edge diagrams $\mathfrak{e} = \varepsilon(\mathfrak{m})$, it is natural to refer to the semicircular arcs as the *edges* of the edge diagram $\mathfrak{e}$ and the small discs as the *vertices* of $\mathfrak{e}$. Those edges whose right end contains a vertex are *vertex-beginner* edges and correspond to canonical tree edges in $\mathfrak{m}$. An edge whose right endpoint intersects an "$\times$" is called a *twisted* edge. If $u = P_i$ and $v = P_k$ are consecutive vertices and $i \leqslant j < k$ where P_j is the endpoint of an edge e, then e and P_j are said to be *incident* to u.

Given an edge diagram $\mathfrak{e}$, its associated topological map can be determined. Form a vertex v homeomorphic to an open disk for each vertex u of $\mathfrak{e}$. Take a series of ribbons, one ribbon for each point P_i incident to u, and join one end of each ribbon to v. For each semicircular arc s of $\mathfrak{e}$ joining points P_i and P_j, join the other ends of the two ribbons corresponding to P_i and P_j. If s is twisted, then a $180°$ twist is introduced to one of the ribbons before they are joined. The surface is completed by sewing discs (faces) to the boundary. The embedded graph is formed in the surface in the obvious manner, by embedding an edge down the middle of the pairs of joined ribbons.

Example of the formation of an edge diagram. Consider the planar embedding of K_4 given in Figure 1. This is done in Figure 6(a). To draw the edge diagram, draw a base line and a number of points labelled $a, A, b, B, \ldots$ where a, A are identified, b, B are identified, ... for this purpose. For each edge e_i of $\mathfrak{m}$, draw a semicircular arc above the base line which joins two points labelled with the position labels that the edge e_i contains, as in Figure 6(b). Place a vertex (black disc) at the left end of the base line, and at the right end of each canonical tree edge (the tree edges are indicated in third diagram of Figure 1: they are the edges on the trace of Algorithm 1). The edge diagram itself does not include the canonical position labels. (These are easily recovered because they appear in order $a, A, b, B, \ldots$.)

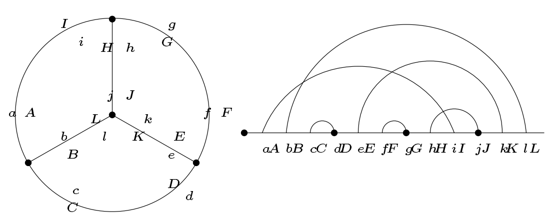

(a) Canonical position labels of $\mathfrak{m}$ (b) Edge diagram of $\mathfrak{m}$ with canonical position labels

Figure 6: Linearising a rooted map $\mathfrak{m}$ into an *edge diagram*

When the rooted map $\mathfrak{m}$ is non-orientable, as in Figure 7(a), there is an additional consideration for each edge. Since $\mathfrak{m}$ is non-orientable, its matching graph is non-bipartite. Then μ_3 sometimes pairs even (uppercase capital letter) position labels with even (uppercase capital letter) position labels. To mark this event, we place an "$\times$" at the right end of the corresponding edge (semicircular arc). This is shown in Figure 7. Note that this cannot occur in a tree edge, so "$\times$" will not be placed on a vertex.

Tracing Boundaries of Faces in an Edge Diagram. In the transformation $\mathfrak{e} = \varepsilon(\mathfrak{m})$, no information is lost ($\mathfrak{m}$ is recoverable from $\mathfrak{e}$). Therefore the face partition $\phi(\mathfrak{m})$ of $\mathfrak{m}$ is directly recoverable from the edge diagram $\mathfrak{e}$. This can be done by tracing along edges. In Figure 8, two faces are traced out in a edge diagram. The rules for tracing the boundary of a face are

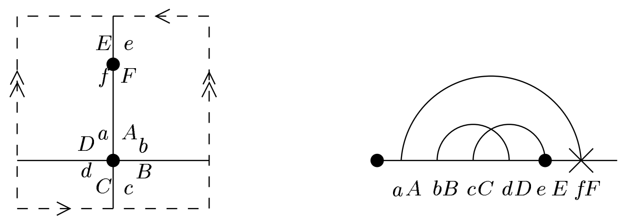

Figure 7: Adding a "×" to the edge diagram of a non-orientable map $\mathfrak{m}$

- trace along a side incident to the face of a semicircular arc (edge),
- if the right end of the semicircular arc has an "×", when the edge has a twist, then the boundary of the face crosses the semicircular arc at the twist,
- trace the boundary of the face along the top side of a segment of the base line except unless the segment of the base line is the first segment, the last segment or a segment incident to a vertex at its right end
- if the trace of the boundary of the face above the base line encounters a segment of the base line which is the first segment, the last segment, or incident with a vertex at its right end, then the trace of the boundary of the face crosses the base line and continues in the direction left or right away from the segment along the bottom side of the base line until it passes underneath another segment of the base line of the same kind and then it crosses the base line to trace the side of another edge,
- the trace of the boundary of the face may pass through itself several times but only finishes once it is again tracing the same side of the edge at which it started.

The canonical edge diagram of a rooted map and the method of tracing the boundaries of faces in an edge diagram are important parts of the construction $\tilde{\tilde{\Xi}}$.

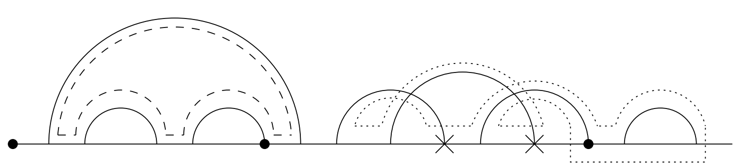

Figure 8: Boundaries of two faces traced in an edge diagram (boundaries of three other faces are not traced)

4 The Quadrangulation Conjecture

Recall that the genus of an orientable map $\mathfrak{m}$ is $g = g(\mathfrak{m}) = \frac{1}{2}(2-\chi)$, where $\chi = f - e + v$ is the Euler characteristic of $\mathfrak{m}$ which has f faces, e edges and v vertices. For any rooted map$\mathfrak{m}$, let $V(\mathfrak{m})$ be its set of vertices. A *dual-strut* of $\mathfrak{m}$ is an edge e of $\mathfrak{m}$ whose associated edge e^∂ in the dual map $\mathfrak{m}^\partial$ is a strut (defined in §3.3). Let $S(\mathfrak{m})$ be the set of dual-struts of $\mathfrak{m}$. Let $\mathcal{O}$ be the set of rooted orientable maps. Let $\mathcal{O}_n$ be the set of maps in $\mathcal{O}$ with n edges.

Definition 3 *For $n \geqslant 1$ and $0 \leqslant g \leqslant \lfloor \frac{n+1}{2} \rfloor$, let*

$$\mathcal{Q}_{g,n} = \{\mathfrak{q} \mid \mathfrak{q} \in \mathcal{O}, g(\mathfrak{q}) = g,\ \mathfrak{q} \textit{ is face-4-regular}\},$$

and

$$\mathcal{A}_{g,n} = \{(\mathfrak{o}, W, F) \mid \mathfrak{o} \in \mathcal{O}_n, W \subseteq V(\mathfrak{o}), F \subseteq S(\mathfrak{o}), \tfrac{|W|}{2} + g(\mathfrak{o}) = g\}.$$

A map $\mathfrak{q} \in \mathcal{Q}_{g,n}$ *is a* quadrangulation *and an object* $\mathfrak{a} \in \mathcal{A}_{g,n}$ *is a* decorated *map.*

A decorated map $\mathfrak{a} = (\mathfrak{o}, W, F)$ may be represented as the orientable map $\mathfrak{o}$ with its vertices in W coloured white, the remaining vertices black, and its dual-struts in $S(\mathfrak{o})$ each bearing a positive or negative *flag*, with F consisting of the dual-struts with negative flags. Table 1(b) gives all the decorated maps in $\mathcal{A}_{1,2}$ using this representation of coloured vertices and flagged dual-struts with $\boxplus$ indicating a positive flag and $\boxminus$ indicating a negative flag. Observe that since g and $g(\mathfrak{o})$ are integers, the number of white vertices is even in any decorated map.

In [4], $\mathcal{A}_{g,n}$ was defined without the use of flagged dual-struts (without any equivalent to F) but rather by counting maps of genus h with multiplicity 4^h. However, since a map of genus h always has $2h$ dual-struts, there are 4^h ways to assign $+/-$ flags to the dual-struts. Both definitions therefore lead to the same cardinality for $\mathcal{A}_{g,n}$.

Character factorisation leads to a proof [4] that:

$$|\mathcal{Q}_{g,n}| = |\mathcal{A}_{g,n}|, \qquad \forall g, n. \tag{2}$$

Let $\mathcal{Q} = \bigcup_{g,n} \mathcal{Q}_{g,n}$ and $\mathcal{A} = \bigcup_{g,n} \mathcal{A}_{g,n}$. Then (2) implies the existence of a bijection $\xi : \mathcal{Q} \to \mathcal{A}$, which preserves the two parameters g and n, that is, $\xi(\mathcal{Q}_{g,n}) = \mathcal{A}_{g,n}$. The authors in [4] proposed the Quadrangulation Conjecture (Conjecture 1) which is that there exists a natural (constructive) bijection, by which is meant a bijection with an algorithmic combinatorial element-wise action, $\Xi : \mathcal{Q} \to \mathcal{A}$, which also preserves g and n.

For such a hypothetical constructive bijection, Ξ, consider a family of more specific bijections $\Xi_{g,n} : \mathcal{Q}_{g,n} \to \mathcal{A}_{g,n}$, defined as the restriction, $\Xi_{g,n} = \Xi|_{\mathcal{Q}_{g,n}}$. In the special instance $g = 0$, prior to [4], Tutte had already proved the special case $|\mathcal{Q}_{0,n}| = |\mathcal{A}_{0,n}|$, and constructed a bijection $M : \bigcup_n \mathcal{Q}_{0,n} \to \bigcup_n \mathcal{A}_{0,n}$, which preserves n. This bijection M is the *medial* construction. Then $M_n = M|_{\mathcal{Q}_{0,n}}$ serves as a valid instance of $\Xi_{0,n}$. But, for $g \geqslant 1$, no general $\Xi_{g,n}$ was known. For $n \leqslant 2$, restrictions of the bijection $\tilde{\tilde{\Xi}}$ give constructions for such bijections. Let $\mathcal{Q}_n = \bigcup_g \mathcal{Q}_{g,n}$ and $\mathcal{A}_n = \bigcup_g \mathcal{A}_{g,n}$. Let $\Xi_n = \Xi|_{\mathcal{Q}_n}$. No general Ξ_n is known that preserves the weight functions for g (since no Ξ is known). However, if the constraint of preserving g is lifted from Ξ_n, then the restriction of $\tilde{\tilde{\Xi}}$ to $\mathcal{Q}_n$ gives a construction for Ξ_n.

5 Constructive Proof of Theorem 1

An overview of the construction of $\tilde{\tilde{\Xi}}$ that proves Theorem 1 is given below. A detailed definition with an example is given in §5.1. Further comments and examples are given in §§5.2–5.4.

An Enumerative Observation. The construction of $\tilde{\tilde{\Xi}}$ is based on the following enumerative observation. By specialising a more general character-based formula it was shown in [4] that

$$|\mathcal{Q}_n| = [x^n]\, 4x \frac{\partial}{\partial x} \log \sum_{n \geqslant 0} \frac{(4n)!}{(2n)!n!16^n} x^n.$$

But it is also true that

$$L_n = |\mathcal{L}_n| = [x^n]\, 4x \frac{\partial}{\partial x} \log \sum_{n \geqslant 0} \frac{(4n)!}{(2n)!n!16^n} x^n, \tag{3}$$

where $\mathcal{L}_n = \{\mathfrak{m} \in \mathcal{L} \mid n(\mathfrak{m}) = n\}$ is the set of rooted, locally orientable maps with n edges. (To show (3), one can use the algebraic method, or use a direct argument with matchings.) Therefore, there exists bijections $\lambda_n : \mathcal{Q}_n \to \mathcal{L}_n$. In fact, we will construct a bijection $\Lambda : \mathcal{Q} \to \mathcal{L}$, which preserves n. Then, we will construct a bijection $A : \mathcal{L} \to \mathcal{A}$, also preserving n, leading to a chain of bijections:

$$\mathcal{Q} \xrightarrow{\ \Lambda\ } \mathcal{L} \xrightarrow{\ A\ } \mathcal{A}$$

and

$$\tilde{\tilde{\Xi}} = A \circ \Lambda. \tag{4}$$

The bijection $\tilde{\tilde{\Xi}} : \mathcal{Q} \to \mathcal{A}$ satisfies

$$\tilde{\tilde{\Xi}}\Big|_{\bigcup_n \mathcal{Q}_{0,n}} = \partial \circ M, \tag{5}$$

where ∂ is the operation of forming the dual, thereby extending Tutte's medial construction, and

$$\tilde{\tilde{\Xi}}(\mathcal{Q}_n) = \mathcal{A}_n, \tag{6}$$

thereby preserving the weight function for n. For $n = 1$ and $n = 2$, one has $g \leqslant \lfloor \frac{3}{2} \rfloor = 1$, so that

$$\begin{aligned} \mathcal{Q}_n &= \mathcal{Q}_{0,n} \cup \mathcal{Q}_{1,n}, \\ \mathcal{A}_n &= \mathcal{A}_{0,n} \cup \mathcal{A}_{1,n}. \end{aligned}$$

Since $\tilde{\tilde{\Xi}}$ is bijective, (5) and (6) imply that $\tilde{\tilde{\Xi}}(\mathcal{Q}_{1,n}) = \mathcal{A}_{1,n}$, and therefore that $\tilde{\tilde{\Xi}}\Big|_{\mathcal{Q}_{1,n}}$ serves as an instance of $\Xi_{1,n}$, for $n \leqslant 2$. (For $n \geqslant 3$, however, our construction $\tilde{\tilde{\Xi}}$ is not guaranteed to yield any more instances of $\Xi_{g,n}$, because $\tilde{\tilde{\Xi}}(\mathcal{Q}_{1,3}) \neq \mathcal{A}_{1,3}$.)

Constituent Bijections of $\tilde{\tilde{\Xi}}$. Recall from (4) that $\tilde{\tilde{\Xi}} = A \circ \Lambda$, for certain bijections $\Lambda : \mathcal{Q} \to \mathcal{L}$ and $A : \mathcal{L} \to \mathcal{A}$. Then (referring to Table 3),

$$\Lambda = R \circ \varepsilon^{-1} \circ \varphi \circ \varepsilon \circ \partial \qquad A = \omega \circ \partial \circ \pi \circ \kappa \circ \varepsilon^{-1} \circ \psi \circ \varepsilon$$

where Table 2 summarises these constituent bijections. Only φ and ψ alter the underlying topological

∂	dualise
ε	form edge diagram
φ	face-bipartise (by adding twists in edge diagram)
R	apply radial construction
ψ	sign and untwist (in edge diagram)
κ	adjust cuts (move their signs to faces)
π	assign parity ($\pm$) to root face
ω	colour vertices (according to $\pm$ signs)

Table 2: The constituent bijections

surface, and only R changes the number of edges (it halves them).

The following is worth noting. Let $\Omega = \varepsilon^{-1} \circ \varphi \circ \varepsilon \circ \partial$. Then $\partial \circ \Omega$ is a bijection between orientable quadrangulations and bipartite locally orientable quadrangulations. Of course, $\Lambda = R \circ \Omega$ and $\tilde{\tilde{\Xi}} = A \circ R \circ \Omega$.

5.1 Details of the Construction of $\tilde{\tilde{\Xi}}$

We now define the construction of $\tilde{\tilde{\Xi}}$, given by (4). This discussion is assisted by working with a particular example, given in Table 3. The table shows, by a sequence of rooted maps and edge diagrams, how to compute the image of $\mathfrak{q}_{1.2.9}$ under $\tilde{\tilde{\Xi}}$. The rooted map (1) in the first stage is the rooted orientable quadrangulation $\mathfrak{q}_{1.2.9} \in \mathcal{Q}_{1,2}$. This map is one of the maps in Table 1(a). The subsequent stages record the progress of the construction $\tilde{\tilde{\Xi}}$. The rooted map (13) of the final stage in the sequence is the decorated rooted orientable map $\mathfrak{a}_{1.2.9} \in \mathcal{A}_{1,2}$.

Transition from one stage to the next is by a bijection and the composition of the constituent bijections gives $\tilde{\tilde{\Xi}} = A \circ R \circ \Omega$. Certain constituent bijections are used more than once. This is indicated by a roman numeral.

	1. $\mathfrak{q}_{1.2.9}$ 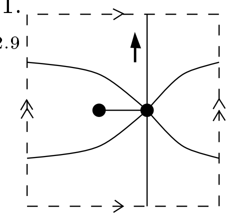	13. $\mathfrak{a}_{1.2.9}$ 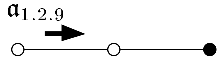	12-13: ω White: $-$ Black: $+$
1-2: ∂ Duality	2. 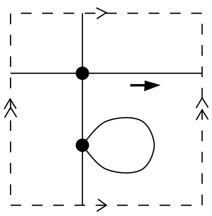	12. 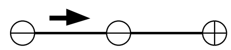	11-12: ∂ Duality
2-3: ε Edge diagram		11. 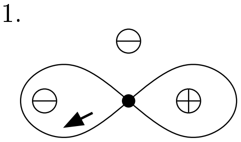	10-11: π Root Face Sign by Parity
	3. 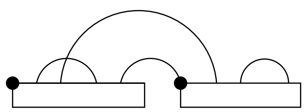	10. 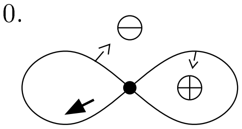	9-10: κ Sign Transfer From Cuts to Associated Faces
3-4: φ Twist by Corner Colours	4. 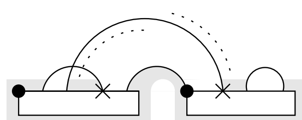	9. 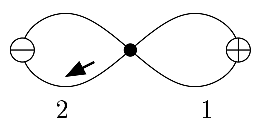	8-9: ε^{-1} Edge diagram
4-5: ε^{-1} Edge diagram	5. 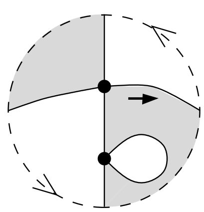	8. 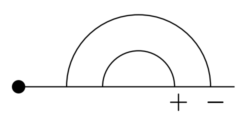	7-8: ψ Twist: $-$
5-6: R Radial Construction	6. 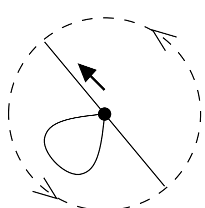	7. 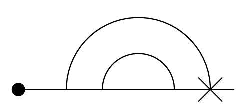	
	6-7: εEdge diagram		

Table 3: Thirteen stages, and the constituent bijections, used in the computation $\widetilde{\Xi}(\mathfrak{q}_{1.2.9}) = \mathfrak{a}_{1.2.9}$.

5.1.1 The Bijection Ω

Stages 1–2: Duality (I). The first constituent bijection is ∂, the operation of taking the dual of a map. This is the first of two instances of duality in the construction of $\tilde{\tilde{\Xi}}$.

The two rooted maps in Stages 1 and 2 of Table 3 are duals of each other. Both maps are in the torus. The vertex and face partition of the map in Stage 1 are [71] and $[4^2]$ and vice versa in the map of Stage 2.

The operation of taking the dual is reversible, because $\partial^{-1} = \partial$. In general, the map in Stage 1 can be any roooted orientable quadrangulation (face-4-regular map) and the map in Stage 2 can be any rooted, orientable, quartic (vertex-4-regular) map.

Stages 2–3: Edge Diagram (I). The second constituent bijection is ε, the operation of forming the edge diagram. This step is the first of four steps in $\tilde{\tilde{\Xi}}$ that convert a map to its edge diagram or vice versa.

Stage 3 of Table 3 is the edge diagram of the Stage 2 map. Recall that the construction of the edge diagram involves an execution of Algorithm 1 on the rooted map.

The constituent function ε, the operation of forming the edge diagram, is a reversible bijection because there is a bijective correspondence between edge diagrams and rooted maps. In general, the edge diagram in Stage 3 may be any diagram that

- has no twists, and
- is quartic in the sense that vertices are located before the first end, at the fifth end, and at every fourth subsequent end, (the ninth, the thirteenth end, and so on),

because the maps in Stage 2 are quartic orientable rooted maps.

Stages 3–4: Adding Twists. The third constituent bijection is φ, the operation of adding twists to an orientable quartic edge diagram in a canonical way such that the result is face-bipartite. To apply the radial construction later in the bijection $\tilde{\tilde{\Xi}}$ it is necessary that the map be face-bipartite. Face-bipartite maps necessarily have all even vertex-degrees, but the converse is not true in general for non-planar maps. The map in Stage 4 is quartic but might not be face biparite. Adding twists does not alter the vertex degrees, so the map in Stage 5 will also be quartic and thus have even vertex degrees. Furthermore, by the following canonical choice of which edges of Stage 4 to twist, Stage 5 is face-bipartite. The canonical tree edges are used to assign colours black and white to the corners of faces locally at vertices. Twists are added to non-tree edges that are inconsistent with respect to the local colouring. The resulting properly face bicoloured map is Stage 5.

The Stage 3 edge diagram in Table 3 represents an orientable, but not face-bipartite, map. We add twists to the edge diagram in such a way that the map becomes face-bipartite. To do this, alternately shade the corners of vertices black and white, and use tree edges to initiate later vertices. It will be easier to visualize this alternating shading procedure if the base line of the edge diagram is replaced by a series of rectangles corresponding to the vertices, as in Figure 9. Then the corners of the edge diagram (which correspond to corners of a map) are more apparent. Begin by shading the corner to which the root belongs. In Figure 9, this first corner extends all the way beneath the first rectangle. On top of the first rectangle, there are three more corners. The outer two of these are left white, while the inner is shaded. Thereby, the corners are alternately shaded and unshaded in the neighbourhood of the first vertex. The shading provides a face-bicolouring, but only for the vicinity of the root vertex. The next step is to alternately shade the corners of the second vertex. Once one corner of the second vertex is shaded, alternation determines the shading of the remaining corners. The first corner of the second vertex may be shaded by using the canonical tree edge joining the second vertex and first edge. Shade the entire length of the tree edge on the side determined by the shaded corner of the first vertex. The shaded side of this tree edge at the second vertex determines its first shaded corner. Once all the corners have been shaded, the resulting shading is a face-bicolouring in the vicinity of the canonical spanning tree. Where the shading may fail to extend as a face-bicolouring is at the non-tree edges. The shaded corners at the ends of a non-tree edge may lie on different sides of the non-tree edge. For these edges, the face-bicoloring fails. To correct this, add a twist in the edge diagram to these failing non-tree edges. In Figure 9, the dashed lines drawn along the sides of non-tree edges begin at the shaded corner of the left (lower) end of the non-tree edge. Note that dashed line of the rightmost non-tree edge arrives at a shaded corner of the right end. Therefore no twist is needed for this non-tree edge. For the two other non-tree edges, however, the dashed lines arrive on the right at unshaded (white) corners. Therefore, twists are added in the edge diagram to these two non-tree edges. The result is the Stage 4 edge diagram.

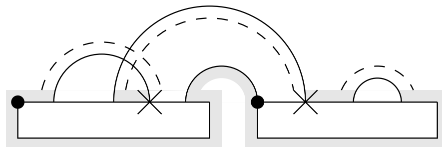

Figure 9: Adding twists to the Stage 3 edge diagram to correct failures to the face-bicolouring

The constituent bijection φ is reversible because its inverse φ^{-1} is the operation of removing all the twists to produce an orientable diagram. In general, the Stage 4 edge diagram is any quartic, face-bipartite, locally orientable edge diagram.

Stages 4–5: Edge Diagram (II). The fourth constituent bijection is ε^{-1}. The rooted map in Stage 5 is the unique locally orientable rooted map whose edge diagram is the one in Stage 4.

In Table 3, the face-colouring of the previous stage is indicated for clarity. Note that the map in Stage 5 of Table 3 is in the projective plane, which is a different surface from the torus, the surface of $\mathfrak{q}_{1.2.9}$.

Clearly, ε^{-1} is reversible, with inverse ε. In general, the maps in Stage 5 can be any vertex-4-regular locally orientable, face-bipartite, rooted map.

5.1.2 The Bijection R

Stages 5–6: Radial Construction. The fifth constituent bijection is R, the radial construction [8]. The radial construction is a dual of Tutte's medial construction. The radial construction R involves constructing a new map from a quartic face-bipartite rooted map by placing a new vertex in each black face and adjoining the new vertices by new edges passing through each quartic vertex. The underlying topological surface of the new map is the same as the original map. The number of edges is halved. The root position of the new map is on the edge that passes through the root vertex of the old map. The root position of the new map is on the same side and end of the edge as the root position of the old map.

The rooted map of the Stage 5 is quartic and face-bipartite (by construction), and the radial construction is applied to construct the map of Stage 5. In Figure 10, the Stage 6 rooted map is superimposed on the Stage 5 rooted map. Stage 6 is represented by a hollow vertex and double lined edges and Stage 5 is represented with solid vertices and single lined edges. Note that the root position in the two rooted maps, each indicated by the tail of an arrow, are such that the arrows can be represented as eminating from the same point.

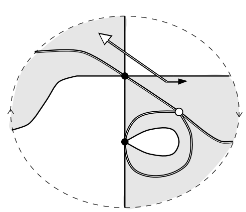

Figure 10: Stages 5 – 6

The radial construction R is reversible. In its inverse construction R^{-1}, a new map is formed by replacing each edge by a vertex and joining these with an edge for corner of the original map. In general, Stage 6 can be any rooted, locally orientable map. (Note that Stage 6 is a halfway point in the computation $\widetilde{\Xi}$ and if Stage 1 is $\mathfrak{q}$ then Stage 6 is $\Lambda(\mathfrak{q})$.)

5.1.3 The Bijection A

Stages 6–7: Edge Diagram (III). The sixth constituent bijection is ε, the operation of forming the edge diagram.

In Table 3, Stage 7 is the edge diagram of Stage 6, which in this instance is a non-orientable map, and consequently the resulting edge diagram in Stage 7 has twists.

As noted previously ε is invertible. In general, Stage 7 of the computation of $\widetilde{\Xi}$ can be any edge diagram.

Stages 7–8: Twists Become Signs. The seventh constituent bijection is ψ, the operation of assigning signs to the some of the arcs (edges) of the edge diagram and removing the twists. The canonical tree edges (arcs) are not assigned a sign. The edges that had twists are assigned a negative sign $(-)$ (such edges will have been canonical non-tree edges). The remaining edges are assigned a postive sign $(+)$.

In Table 3, there is only one vertex, so there are no canonical tree edges. The twists are removed from the edge diagram of Stage 7 to yield an orientable edge diagram for the eighth stage. Formerly twisted edges are marked by a negative sign $(-)$, while untwisted non-tree edges are marked by a positive sign $(+)$.

The operation ψ is reversible because the information about the location of the twists that are removed is retained by the location of the negative signs. In other words, ψ^{-1} is computed by removing the signs and adding twists to the edges that had a negative sign. In general, the structures in Stage 8 of the computation of $\widetilde{\Xi}$ are edge diagrams with signed canonical non-tree edges.

Stages 8–9: Edge Diagram (IV). The eighth constituent bijection is essentially ε^{-1}, the operation of recovering a rooted map from its edge diagram, but in addition the signs of edges in the edge diagram are transferred to the corresponding edges of the rooted map.

In Table 3, both edges are non-tree edges. In Figure 11, on the left is the Stage 8 edge diagram with its canonical labelling $(a, A, b, B, \dots)$, and on the right is the Stage 9 rooted map with its canonical labeling. Note that negative sign is on the edge with labels $aAdD$ in both stages.

Figure 11: Determining the map (Stage 9) from the edge diagram (Stage 8), and carrying over the edge-signs

The constituent bijection ε^{-1} is reversible because, as before, of the correspondence between rooted maps and edges diagrams and also because, considering the signs of edges, for each edge diagram and associated rooted map there is a correspondence between the edges of one and the other. In general, the rooted maps in Stage 9 of the computation of $\widetilde{\Xi}$ are orientable maps with signed canonical non-tree edges.

Stages 9–10: Sign the Non-root Faces. The ninth constituent bijection is κ, the operation of transferring signs from cuts to their associated non-root faces. The result of the operation κ is that the rooted map has $\pm$ signs associated exactly with its non-root faces and its struts.

In Table 3, the rooted map of Stage 9 has no tree edges and has no struts, that is, both its edges are cuts. Recall from §3.3 on page 7 the canonical association between cuts and non-root faces, which is given in Figure 12(b). Each non-root face is assigned a sign equal to the sign that was assigned to the associated cut.

The constituent bijection κ is reversible because, to reverse it, the signs of non-root faces are simply transferred to their associated cuts. In general, Stage 10 of the computation of $\widetilde{\Xi}$ is a rooted map with signed non-root faces and struts.

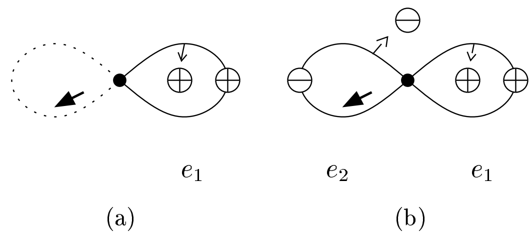

Figure 12: Transferring signs from the cuts of Stage 9 to the faces of Stage 10

Stages 10–11: Sign the Root Face. The tenth constituent bijection is π, which is the operation of assigning the unique $\pm$ sign to the root face that makes the parity of the total number of faces with a negative sign even.

Stage 10 of Table 3 has an odd number of negatively signed faces. Therefore Stage 11 of Table 3 has a negatively signed root face.

The constituent bijection π is reversible because the operation to reverse it is remove the sign of the root face. In general, Stage 11 of the computation of $\widetilde{\Xi}$ is an orientable rooted map with signed faces, an even number of which are negative, and signed struts.

Stages 11–12: Duality (ii). The eleventh constituent bijection of $\widetilde{\widetilde{\Xi}}$ is the operation ∂ of forming the dual of a rooted map, where in addition the signed subobjects in the rooted map are transferrred to the associated dual subojbects in the dual. That is, the $\pm$ sign of an edge that is a strut is transferred to the associated dual edge (which is a dual-strut) and the $\pm$ sign of a face is transferred to the associated vertex which is dual to the face.

Stage 12 of Table 3 has the dual rooted map of the rooted map in Stage 11 and the signs are the vertices are carried over from the signs of the faces. In this instance, Stage 12 has no dual-struts, so the edges are unsigned.

This constituent bijection ∂ is reversible because the signs can transferred back uniquely, and duality is an involution of rooted maps. In general, Stage 12 of the computation $\widetilde{\widetilde{\Xi}}$ can be any orientable rooted map with signed vertices, an even number of which are negative, and signed dual-struts.

Stages 12–13: Signs to Colours. The twelfth and last constituent bijection is ω, the operation of converting the negative vertices to white-coloured vertices, and positive vertices to black-coloured vertices. The signs of the dual-struts are not changed. The result of ω is a decorated map according to §4.

Stage 13 of Table 3 is result of ω applied to Stage 12, and is also the decorated map $\mathfrak{a}_{1.2.9}$ of Table 1(b).

The operation ω is reversible because it can be reversed by replacing white vertices by negative vertices and black vertices by positive vertices. In general, Stage 13 of the computation of $\widetilde{\widetilde{\Xi}}$ can be any decorated map as described in §4, which is orientable rooted map, with black and white vertices, an even number of which are white, and signed dual-struts.

5.2 Preservation of $g = 0$ by $\widetilde{\widetilde{\Xi}}$

Let $\mathfrak{q} \in \mathfrak{Q}$ be planar. Simplify $\Lambda(\mathfrak{q})$ as follows. Since $\partial(\mathfrak{q})$ is planar and quartic, it follows (by a generalisation of the Jordan Curve Theorem) that $\partial(\mathfrak{q})$ is face-bipartite. Therefore φ does not twist any edges of the edge diagram of $\partial(\mathfrak{q})$. Hence $\Lambda(\mathfrak{q}) = R \circ \partial(\mathfrak{q}) = M(\mathfrak{q})$. It follows that $\Lambda(\mathfrak{q})$ is planar. In particular $\Lambda(\mathfrak{q})$ is orientable so, more generally, let $\mathfrak{o}$ be some orientable rooted map, and consider $A(\mathfrak{o})$. Since $\mathfrak{o}$ is orientable, its edge diagram has no twists, and thus ψ has no effect on the edge diagram other than to assign $+$ to each edge. Since all the edge signs are $+$, then κ makes all the non-root face signs $+$, and π then makes the root face sign $+$. At this point (Stage 11), there is a rooted orientable map with all positive faces and all positive struts. Applying ∂ leaves all the signs positive, and leaves the map orientable. Then ω produces no white vertices, because there are no negative signs.

The final result, $\mathfrak{a} = A(\mathfrak{o})$, has black vertices and positive dual-struts. Thus $\mathfrak{a} = (\mathfrak{o}', \varnothing, \varnothing)$. In fact, $\mathfrak{o}' = \partial(\mathfrak{o})$, because none of the operations κ, π and ω affects the underlying map, while ψ does not affect the underlying map because there was nothing to untwist in the edge diagram of $\mathfrak{o}$. Thus $A(\mathfrak{o}) = (\partial(\mathfrak{o}), \varnothing, \varnothing)$. In particular, when $\mathfrak{q}$ is planar, then $\widetilde{\Xi}(\mathfrak{q}) = A \circ \Lambda(\mathfrak{q}) = (\partial \circ \Lambda(\mathfrak{q}), \varnothing, \varnothing) = (\partial \circ M(\mathfrak{q}), \varnothing, \varnothing)$. Therefore $\widetilde{\Xi}$ extends Tutte's medial construction.

5.3 Action of $\widetilde{\Xi}$ on $\mathcal{Q}_{1,2}$

Table 1(a) shows the set $\mathcal{Q}_{1,2}$, with labels $\mathfrak{q}_{1.2.i}$ for $i = 1, \ldots, 15$. These are all orientable rooted quadrangulations with two vertices on the torus. Table 1(b) shows the members of the set $\mathcal{A}_{1,2}$, with labels $\mathfrak{a}_{1.2.i}$ for $i = 1, \ldots, 15$, such that $\mathfrak{a}_{1.2.i} = \widetilde{\Xi}\mathfrak{q}_{1.2.i}$. For each of these rooted maps, some even number $2w$ of the vertices may be coloured white. If h is the genus of one of these maps, then $g = w + h = 1$. (This is part of the definition of $\mathcal{A}_{1,2}$.) The definition of $\widetilde{\Xi}$ describes how to determine the rooted map, the black and white vertex colours, and the flagged dual-struts of a decorated map $\mathfrak{a} = \widetilde{\Xi}(\mathfrak{q})$ from any given quadrangulation $\mathfrak{q}$. Thus, although perhaps not the simplest construction, it provides independent proof that $|\mathcal{Q}_{1,2}| = |\mathcal{A}_{1,2}|$, altogether distinct from (2) and from exhaustive listing of both sets.

5.4 Comments on Genus Preservation

When $n = 3$, the parameter $g \in \{0, 1, 2\}$. As we have shown, $\widetilde{\Xi}$ preserves the parameter when $g = 0$. But there are quadrangulations $\mathfrak{q} \in \mathcal{Q}_{1,3}$ such that $\widetilde{\Xi}(\mathfrak{q}) \in \mathcal{A}_{2,3}$, such as the quadrangulation in Stage 1 of Table 4. Therefore, for $n = 3$, the parameter g is not preserved by $\widetilde{\Xi}$ unless it is zero.

Observe that $|\mathcal{Q}_3| = |\mathcal{L}_3| = |\mathcal{A}_3| = 297$, and that $\widetilde{\Xi}$ puts these 297 objects in bijection with each other. The pre-image and image sets, and their sizes, which are to be preserved by Ξ, are

$$|\mathcal{Q}_{0,3}| = |\mathcal{A}_{0,3}| = 54, \qquad |\mathcal{Q}_{1,3}| = |\mathcal{A}_{1,3}| = 198, \qquad |\mathcal{Q}_{2,3}| = |\mathcal{A}_{2,3}| = 45.$$

Then we may note $\widetilde{\Xi}$ provides a bijection between $\mathcal{Q}_{1,3} \cup \mathcal{Q}_{2,3}$ and $\mathcal{A}_{1,3} \cup \mathcal{A}_{2,3}$, which are sets of cardinality $243 = 198 + 45$. But $\widetilde{\Xi}$ combines up some of the 198 elements of one set with the 45 elements of the other set. However, Ξ must not separate these elements.

It is possible that $\widetilde{\Xi}$ can be corrected to become the bijection Ξ that we seek. The case of $n = 3$ is first case where $\widetilde{\Xi}$ fails to preserve genus, and therefore should be the starting point of an investigation of how to correct $\widetilde{\Xi}$. For example, consider bijections $\widehat{\Xi} = A \circ \Pi \circ \Lambda$, where $\Pi : \mathcal{L} \to \mathcal{L}$ is a bijection of locally orientable, rooted maps. If Π preserves n (the number of edges) and preserves planarity, then $\widehat{\Xi}$ will share the properties of $\widetilde{\Xi}$: it preserves n and $g = 0$. However, it may be that there is a particular choice of Π that causes $\widehat{\Xi}$ to preserve all values of g. This observation applies to every stage of the computation $\widetilde{\Xi}$: one may insert a permutation amongst the class of rooted maps (or edge diagrams) of that stage. For example, it may also apply to the particular choice of canonical labels or spanning tree.

References

[1] Brown, D. R. L. *Differential Equations and Depth First Search for Enumeration of Maps in Surfaces.* PhD thesis, University of Waterloo, 2000.

[2] Goulden, I. P., and Jackson, D. M. Connection coefficients, matchings, maps and combinatorial conjectures for Jack symmetric functions. *Trans. Amer. Math. Soc. 348* (1996), 873–892.

[3] Jackson, D. M., Perry, M. J., and Visentin, T. I. Factorisations for partition functions of random Hermitian matrix models. *Comm. Math. Phys. 179* (1996), 25–59.

[4] Jackson, D. M., and Visentin, T. I. A character-theoretic approach to embeddings of rooted maps in an orientable surface of given genus. *Trans. Amer. Math. Soc. 322* (1990), 343–363.

[5] Lehman, A. B., and Walsh, T. R. S. Counting rooted maps by genus. I. *J. Combin. Theory, Ser. B 13* (1972), 192–218.

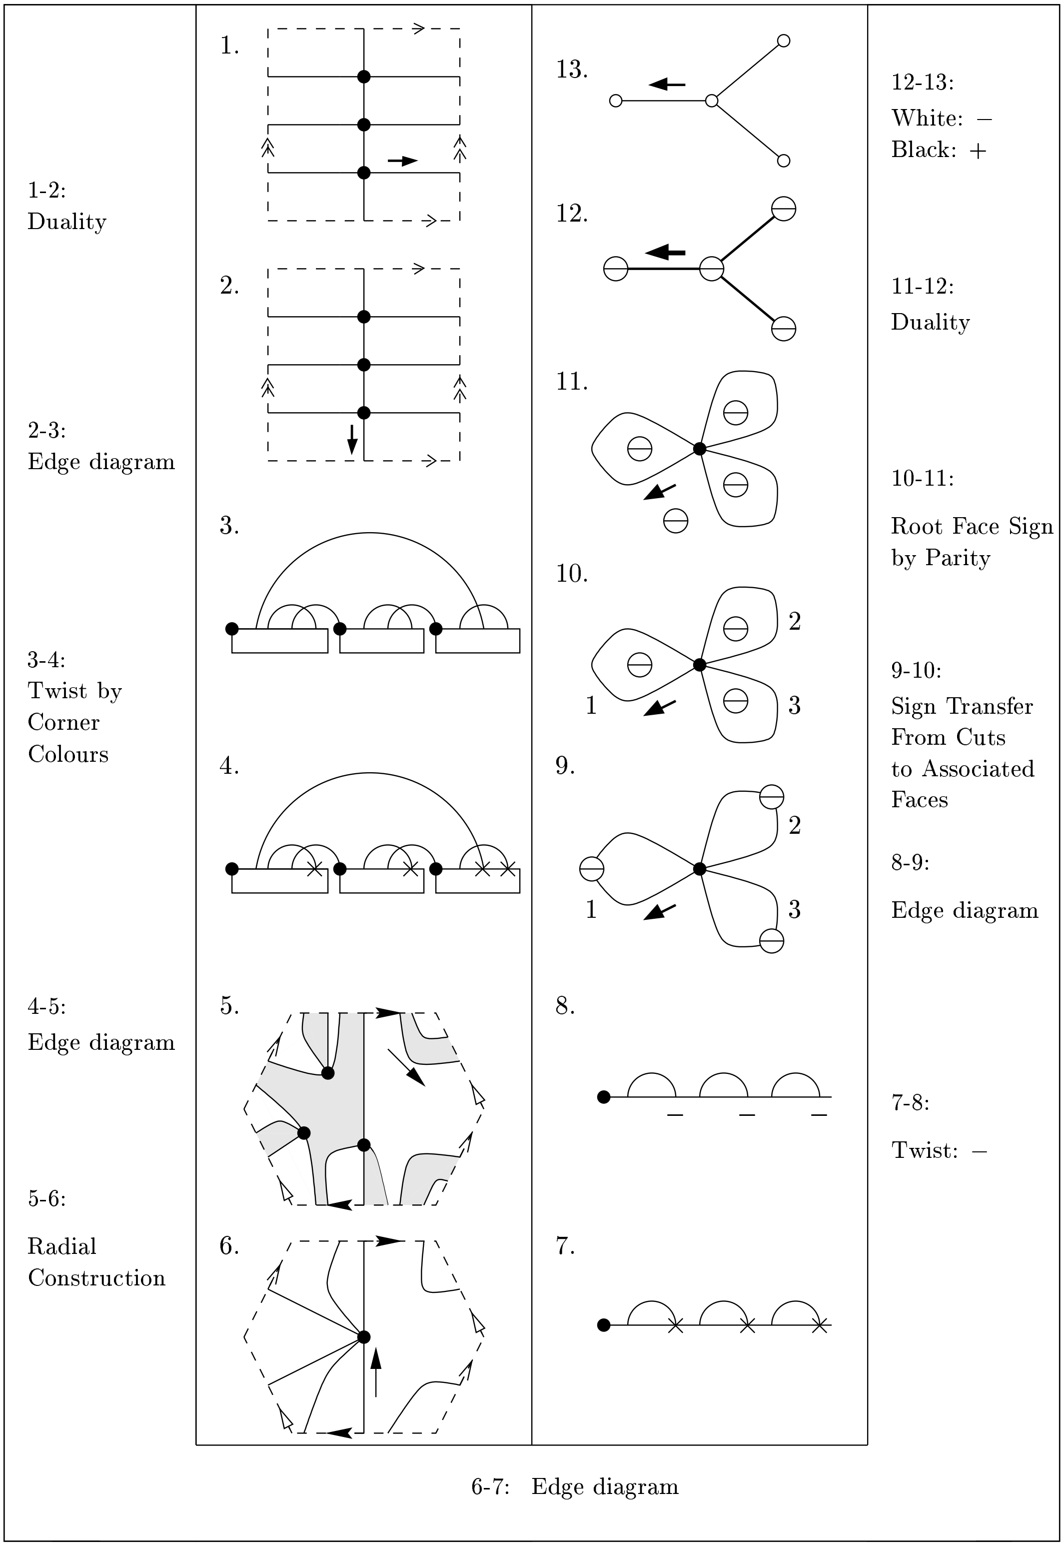

Table 4: A counterexample to preservation of g. ($g = 1$ in Stage 1 and $g = 2$ in Stage 13)

[6] Lehman, A. B., and Walsh, T. R. S. Counting rooted maps by genus. II. *J. Combin. Theory, Ser. B 13* (1972), 122–141.

[7] Penner, R. C. Perturbative series and the moduli space of Riemann surfaces. *J. Differential Geometry 27* (1988), 35–53.

[8] Schaeffer, G. *Conjugaison d'arbres et cartes combinatoires aléatoires.* PhD thesis, L'Université Bordeaux I, 1998.

[9] Tutte, W. T. A census of planar maps. *Canad. J. Math. 15* (1963), 249–271.

[10] Tutte, W. T. *Graph theory*, vol. 21 of *Encyclopedia of Math. and its Applications.* Addison-Wesley Publishing Co., Reading, Mass., 1984.

[11] Walsh, T. R. S. *Combinatorial Enumeration of Non-Planar Maps.* PhD thesis, University of Toronto, 1971.